
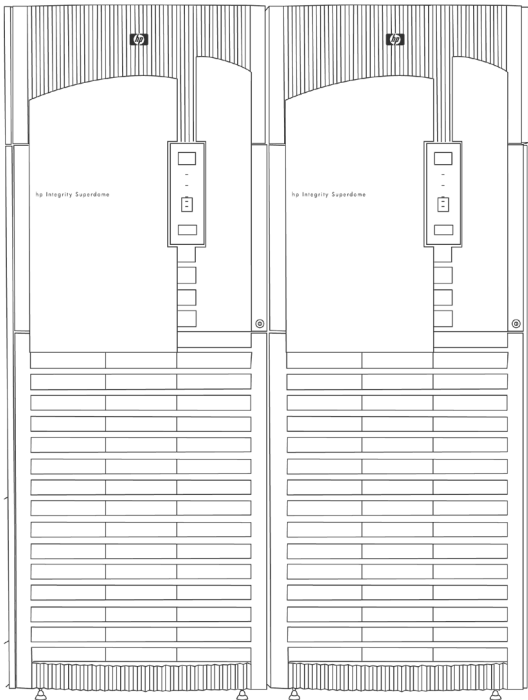


Overview



Overview

At A Glance

The latest release of Superdome, HP Integrity Superdome supports the new and improved sx1000 chip set. HP Integrity Superdome supports Itanium 2 1.5-GHz processors in mid 2003, and the next generation PA RISC processor, PA 8800 and the mx2 processor module based on two Itanium 2 processors in early 2004.

Throughout the rest of this document, the term HP Integrity Superdome with Itanium 2 1.5-GHz processors will be referred to as simply "Superdome".

Superdome with Itanium 2 1.5-GHz processors showcases HP's commitment to delivering a 64-way Itanium server and superior investment protection. It is the dawn of a new era in high end computing with the emergence of commodity based hardware.

Superdome supports a multi OS environment. The multi OS environment offered by Superdome is listed below.

NOTE: Superdome supports both Red Hat Enterprise Linux AS 3 and Debian Linux. Throughout the rest of this document, the two flavors will be collectively referred to as Linux.

NOTE: For information on upgrades from existing Superdome systems to HP Integrity Superdome systems, please refer to the "Upgrade" section.

NOTE: This information can also be found in ESP at:

<http://esp.mayfield.hp.com:2000/nav24/ppos/servers/gen/PriceAvailConfig/59815790/cgch2/cgch2sub8>

HP-UX 11i version 2

- Improved performance over PA 8700
- Investment protection through upgrades from existing Superdomes to next generation Itanium 2 processors

Windows Server 2003 Datacenter Edition for Itanium 2

- Extension of industry standard computing with Windows further into the enterprise data center
- Increased performance and scalability over 32-bit implementations
- Lower cost of ownership versus proprietary solutions
- Ideal for scale up database opportunities such as SQL Server 2000 (64-bit)
- Ideal for database consolidation opportunities such as consolidation of legacy 32-bit versions of SQL Server 2000 to SQL Server 2000 (64-bit)

Linux

- Extension of industry standard computing with Linux further into the enterprise data center
 - Lower cost of ownership versus proprietary solutions
-

Overview

Superdome Service Solutions

Superdome continues to provide the same positive Total Customer Experience via industry-leading HP Services, as with existing Superdome servers. The HP Services component of Superdome is as follows:

- HP customers have consistently achieved higher levels of satisfaction when key components of their IT infrastructures are implemented using the **Solution Life Cycle**. The Solution Life Cycle focuses on rapid productivity and maximum availability by examining customers' specific needs at each of five distinct phases (plan, design, integrate, install, and manage) and then designing their Superdome solution around those needs. HP offers three pre configured service solutions for Superdome that provides customers with a choice of lifecycle services to address their own individual business requirements.
 - **Foundation Service Solution:** This solution reduces design problems, speeds time to production, and lays the groundwork for long term system reliability by combining pre installation preparation and integration services, hands on training and reactive support. This solution includes HP Support Plus 24 to provide an integrated set of 24x7 hardware and software services as well as software updates for selected HP and third party products.
 - **Proactive Service Solution:** This solution builds on the Foundation Service Solution by enhancing the management phase of the Solution Life Cycle with HP Proactive 24 to complement your internal IT resources with proactive assistance and reactive support. Proactive Service Solution helps reduce design problems, speed time to production, and lay the groundwork for long term system reliability by combining pre installation preparation and integration services with hands on staff training and transition assistance. With HP Proactive 24 included in your solution, you optimize the effectiveness of your IT environment with access to an HP certified team of experts that can help you identify potential areas of improvement in key IT processes and implement necessary changes to increase availability.
 - **Critical Service Solution:** Mission Critical environments are maintained by combining proactive and reactive support services to ensure maximum IT availability and performance for companies that can't tolerate downtime without serious business impact. Critical Service Solution encompasses the full spectrum of deliverables across the Solution Lifecycle and is enhanced by HP Critical Service as the core of the management phase. This total solution provides maximum system availability and reduces design problems, speeds time to production, and lays the groundwork for long term system reliability by combining pre installation preparation and integration services, hands on training, transition assistance, remote monitoring, and mission critical support. As part of HP Critical Service, you get the services of a team of HP certified experts that will assist with the transition process, teach your staff how to optimize system performance, and monitor your system closely so potential problems are identified before they can affect availability.
- **HP's Mission Critical Partnership:** This service offering provides customers the opportunity to create a custom agreement with Hewlett Packard to achieve the level of service that you need to meet your business requirements. This level of service can help you reduce the business risk of a complex IT infrastructure, by helping you align IT service delivery to your business objectives, enable a high rate of business change, and continuously improve service levels. HP will work with you proactively to eliminate downtime, and improve IT management processes. S
- **Service Solution Enhancements:** HP's full portfolio of services is available to enhance your Superdome Service Solution in order to address your specific business needs. Services focused across multi operating systems as well as other platforms such as storage and networks can be combined to compliment your total solution.

Standard Features

System	HP UX 11i version 2		Windows Server 2003 Datacenter edition		Linux	
	Minimum	Maximum (in one partition)	Minimum	Maximum (in one partition)	Minimum	Maximum (in one partition)
Superdome 16-way	<ul style="list-style-type: none"> ● 2 CPUs ● 2 GB Memory ● 1 Cell Board ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 16 CPUs ● 128 GB Memory ● 4 Cell Boards ● 4 PCI X Chassis ● 4 npars max 	<ul style="list-style-type: none"> ● 2 CPUs ● 2 GB Memory ● 1 Cell Board ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 16 CPUs ● 128 GB Memory ● 4 Cell Boards ● 4 PCI X Chassis ● 4 npars max 	<ul style="list-style-type: none"> ● 2 CPUs ● 2 GB Memory ● 1 Cell Board ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 8 CPUs ● 64 GB Memory ● 2 Cell Boards ● 1 PCI X Chassis ● 4 npars max
Superdome 32-way	<ul style="list-style-type: none"> ● 2 CPUs ● 2 GB Memory ● 1 Cell Board ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 32 CPUs ● 256 GB Memory ● 8 Cell Boards ● 8 PCI X Chassis ● 8 npars max ● IOX required if more than 4 npars. 	<ul style="list-style-type: none"> ● 2 CPUs ● 2 GB Memory ● 1 Cell Board ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 32 CPUs ● 256 GB Memory ● 8 Cell Boards ● 8 PCI X Chassis ● 8 npars max ● IOX required if more than 4 npars. 	<ul style="list-style-type: none"> ● 2 CPUs ● 2 GB Memory ● 1 Cell Board ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 8 CPUs ● 64 GB Memory ● 2 Cell Boards ● 1 PCI X Chassis ● 8 npars max ● IOX required if more than 4 npars
Superdome 64-way	<ul style="list-style-type: none"> ● 6 CPUs ● 6 GB memory ● 3 Cell Boards ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 64 CPUs ● 512 GB Memory ● 16 Cell Boards ● 16 PCI X Chassis ● 16 npars max ● IOX required if more than 8 npars. 	<ul style="list-style-type: none"> ● 6 CPUs ● 6 GB memory ● 3 Cell Boards ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 64 CPUs ● 512 GB Memory ● 16 Cell Boards ● 16 PCI X Chassis ● 16 npars max ● IOX required if more than 8 npars. 	<ul style="list-style-type: none"> ● 6 CPUs ● 6 GB memory ● 3 Cell Boards ● 1 PCI X Chassis 	<ul style="list-style-type: none"> ● 8 CPUs ● 64 GB Memory ● 2 Cell Boards ● 1 PCI X Chassis ● 16 npars max ● IOX required if more than 8 npars.
Standard Hardware Features	<ul style="list-style-type: none"> ● Redundant Power supplies ● Redundant Fans ● Factory integration of memory and I/O cards ● Installation Guide, Operator's Guide and Architecture Manual ● HP site planning and installation ● One year warranty with same business day on site service response 					

Configuration

There are three basic building blocks in the Superdome system architecture: the cell, the crossbar backplane and the PCI X based I/O subsystem.

Cabinets

Starting with the sx1000 chip set, Superdome servers will be released with the Graphite color. A Superdome system will consist of up to four different types of cabinet assemblies:

- One Superdome left cabinet.
- No more than one Superdome right cabinet (only Superdome 64-way system)
- The Superdome cabinets contain all of the processors, memory and core devices of the system. They will also house most (usually all) of the system's PCI X cards. Systems may include both left and right cabinet assemblies containing, a left or right backplane respectively.
- One or more HP Rack System/E cabinets. These 19-inch rack cabinets are used to hold the system peripheral devices such as disk drives.
- Optionally, one or more I/O expansion cabinets (Rack System/E). An I/O expansion cabinet is required when a customer requires more PCI X cards than can be accommodated in their Superdome cabinets.

Superdome cabinets will be serviced from the front and rear of the cabinet only. This will enable customers to arrange the cabinets of their Superdome system in the traditional row fashion found in most computer rooms. The width of the cabinet will accommodate moving it through common doorways in the U.S. and Europe. The intake air to the main (cell) card cage will be filtered. This filter will be removable for cleaning/replacement while the system is fully operational.

A status display will be located on the outside of the front and rear doors of each cabinet. The customer and field engineers can therefore determine basic status of each cabinet without opening any cabinet doors.

Superdome 16-way and Superdome 32-way systems are available in single cabinets. Superdome 64 way systems are available in dual cabinets.

Each cabinet may contain a specific number of cell boards (consisting of CPUs and memory) and I/O. See the following sections for configuration rules pertaining to each cabinet.

Cells (CPUs and Memory)

A cell, or cell board, is the basic building block of a Superdome system. It is a symmetric multi processor (SMP), containing up to 4 processor modules and up to 16 GB of main memory using 512 MB DIMMs or up to 32 GB of main memory using 1 GB DIMMs. It is also possible to mix 512 MB and 1 GB DIMMs on the same cell board. A connection to a 12 slot PCI X card cage is optional for each cell.

The Superdome cell boards shipped from the factory are offered with 2 processors or 4 processors. These cell boards are different from those that were used in the previous releases of Superdome.

The cell boards can contain a minimum of 2 (for 2-way cell boards) and 4 (for 4-way cell boards) active processors.

The Superdome cell board contains:

- Itanium 2 1.5-GHz CPUs (up to 4 processor modules)
- Cell controller ASIC (application specific integrated circuit)
- Main memory DIMMs (up to 32 DIMMs per board in 4 DIMM increments, using 512 MB or 1 GB DIMMs - or some combination of both.)
- Voltage Regulator Modules (VRM)
- Data buses
- Optional link to 12 PCI X I/O slots

Crossbar Backplane

Each crossbar backplane contains two sets of two crossbar chips that provide a non blocking connection between eight cells and the other backplane. Each backplane cabinet can support up to eight cells or 32 processors (in a Superdome 32-way in a single cabinet). A backplane supporting four cells or 16 processors would result in a Superdome 16-way. Two backplanes can be linked together with flex cables to produce a cabinet that can support up to 16 cells or 64 processors (Superdome 64-way in dual cabinets).

Configuration

I/O Subsystem

Each I/O chassis provides twelve PCI X slots; eight standard and four high bandwidth PCI X slots. There are two I/O chassis in an I/O Chassis Enclosure (ICE). Each I/O chassis connects to one cell board and the number of I/O chassis supported is dependent on the number of cells present in the system. If a PCI card is inserted into a PCI X slot, the card cannot take advantage of the faster slot.

Each Superdome cabinet supports a maximum of four I/O chassis. The optional I/O expansion cabinet can support up to six I/O chassis.

A 4 cell Superdome (16-way) supports up to four I/O chassis for a maximum of 48 PCI X slots.

An 8 cell Superdome supports up to eight I/O chassis for a maximum of 96 PCI X slots. Four of these I/O chassis will reside in an I/O expansion cabinet.

A 16-cell Superdome supports up to sixteen I/O chassis for a maximum of 192 PCI X slots. Eight of these I/O chassis will reside in two I/O expansion cabinets (either six chassis in one I/O expansion cabinet and two chassis in the other, or four chassis in each).

Core I/O

The core I/O in Superdome provides the base set of I/O functions required by every Superdome partition. Each partition must have at least one core I/O card in order to boot. Multiple core I/O cards may be present within a partition (one core I/O card is supported per I/O backplane); however, only one may be active at a time. Core I/O will utilize the standard long card PCI X form factor but will add a second card cage connection to the I/O backplane for additional non PCI X signals (USB and utilities). This secondary connector will not impede the ability to support standard PCI X cards in the core slot when a core I/O card is not installed.

Any I/O chassis can support a Core I/O card that is required for each independent partition. A system configured with 16 cells, each with its own I/O chassis and core I/O card could support up to 16 independent partitions. Note that cells can be configured without I/O chassis attached, but I/O chassis cannot be configured in the system unless attached to a cell.

HP-UX Core I/O (A6865A)

The core I/O card's primary functions are:

- Partitions (console support) including USB and RS 232 connections
- 10/100Base T LAN (general purpose)

Other common functions, such as Ultra/Ultra2 SCSI, Fibre Channel, and Gigabit Ethernet, are not included on the core I/O card. These functions are, of course, supported as normal PCI X add in cards.

The unified 100Base T Core LAN driver code searches to verify whether there is a cable connection on an RJ 45 port or on an AUI port. If no cable connection is found on the RJ 45 port, there is a busy wait pause of 150 ms when checking for an AUI connection. By installing the loopback connector (description below) in the RJ 45 port, the driver would think an RJ 45 cable was connected and would not continue to search for an AUI connection, hence eliminate the 150 ms busy wait state:

Product/ Option Number	Description
A7108A	RJ 45 Loopback Connector
0D1	Factory integration RJ 45 Loopback Connector

Windows Core I/O (A6865A and A7061A and optional VGA/USB A6869A)

For Windows Server 2003, two core I/O cards are required: the Superdome core I/O card (A6865A) and a 1000Base T LAN card (A7061A). The use of Graphics/USB card (A6869A) is optional and not required.

Configuration

I/O Expansion Cabinet	The I/O expansion functionality is physically partitioned into four rack mounted chassis-the I/O expansion utilities chassis (XUC), the I/O expansion rear display module (RDM), the I/O expansion power chassis (XPC) and the I/O chassis enclosure (ICE). Each ICE supports up to two 12-slot PCI-X chassis.
Field Racking	<p>The only field rackable I/O expansion components are the ICE and the 12-slot I/O chassis. Either component would be field installed when the customer has ordered additional I/O capability for a previously installed I/O expansion cabinet.</p> <p>No I/O expansion cabinet components will be delivered to be field installed in a customer's existing rack other than a previously installed I/O expansion cabinet. The I/O expansion components were not designed to be installed in racks other than Rack System E. In other words, they are not designed for Rosebowl I, pre merger Compaq, Rittal, or other third party racks.</p> <p>The I/O expansion cabinet is based on a modified HP Rack System E and all expansion components mount in the rack. Each component is designed to install independently in the rack. The Rack System E cabinet has been modified to allow I/O interface cables to route between the ICE and cell boards in the Superdome cabinet. I/O expansion components are not designed for installation behind a rack front door. The components are designed for use with the standard Rack System E perforated rear door.</p>
I/O Chassis Enclosure (ICE)	<p>The I/O chassis enclosure (ICE) provides expanded I/O capability for Superdome. Each ICE supports up to 24 PCI-X slots by using two 12-slot Superdome I/O chassis. The I/O chassis installation in the ICE puts the PCI-X cards in a horizontal position. An ICE supports one or two 12-slot I/O chassis. The I/O chassis enclosure (ICE) is designed to mount in a Rack System E rack and consumes 9U of vertical rack space.</p> <p>To provide online addition/replacement/deletion access to PCI or PCI-X cards and hot swap access for I/O fans, all I/O chassis are mounted on a sliding shelf inside the ICE.</p> <p>Four (N+1) I/O fans mounted in the rear of the ICE provide cooling for the chassis. Air is pulled through the front as well as the I/O chassis lid (on the side of the ICE) and exhausted out the rear. The I/O fan assembly is hot swappable. An LED on each I/O fan assembly indicates that the fan is operating.</p>
Cabinet Height and Configuration Limitations	Although the individual I/O expansion cabinet components are designed for installation in any Rack System E cabinet, rack size limitations have been agreed upon. IOX Cabinets will ship in either the 1.6 meter (33U) or 1.96 meter (41U) cabinet. In order to allay service access concerns, the factory will not install IOX components higher than 1.6 meters from the floor. Open space in an IOX cabinet will be available for peripheral installation.
Peripheral Support	All peripherals qualified for use with Superdome and/or for use in a Rack System E are supported in the I/O expansion cabinet as long as there is available space. Peripherals not connected to or associated with the Superdome system to which the I/O expansion cabinet is attached may be installed in the I/O expansion cabinet.
Server Support	<p>No servers except those required for Superdome system management such as Superdome Support Management Station or ISEE may be installed in an I/O expansion.</p> <p>Peripherals installed in the I/O expansion cabinet cannot be powered by the XPC. Provisions for peripheral AC power must be provided by a PDU or other means.</p>
Standalone I/O Expansion Cabinet	If an I/O expansion cabinet is ordered alone, its field installation can be ordered via option 750 in the ordering guide (option 950 for Platinum Channel partners).

Configuration

DVD Solution

The DVD solution for Superdome requires the following components. These components are recommended per partition, although it is acceptable to have only one DVD solution and connect it to one partition at a time. External racks A4901A and A4902A must also be ordered with the DVD solution.

Superdome DVD Solutions

Description	Part Number	Option Number
PCI Ultra160 SCSI Adapter or PCI-X Dual channel Ultra160 SCSI Adapter	A6828A or A6829A	0D1
PCI Ultra160 SCSI Adapter or PCI-X Dual channel Ultra 160 SCSI Adapter (Windows Server 2003)	A7059A or A7060A	0D1
Surestore Tape Array 5300	C7508AZ	
DVD (recommend one per partition)	C7499A	0D1
DDS-4 (opt.)/DAT40 (DDS-5/DAT 72 is also supported. Product number is Q1524A)	C7497A	0D1
Jumper SCSI Cable for DDS-4 (optional) ¹	C2978B	0D1
SCSI Cable 2,3	C2363B ²	0D1
SCSI Terminator	C2364A	0D1

1. 0.5-meter HD HDT568 is required if DDS-4 is used.
2. 5-meter multi-mode VH-HD68TS available now (C2365B #0D1) and can be used in place of the 10- meter cable on solutions that will be physically compact
3. 10-meter multi-mode VH-HD68TS available now (C2363B, #0D1)

Partitions

Superdome can be configured with hardware partitions, (npars). Given that HP-UX 11i version 2 does not support virtual partitions (vpars), Superdome systems running HP-UX 11i version 2 do not support vpars.

A hardware partition (npar) consists of one or more cells that communicate coherently over a high bandwidth, low latency crossbar fabric. Individual processors on a single-cell board cannot be separately partitioned. Hardware partitions are logically isolated from each other such that transactions in one partition are not visible to the other hardware partitions within the same complex.

Each npar runs its own independent operating system. Different npars may be executing the same or different revisions of an operating system, or they may be executing different operating systems altogether. Superdome supports HP-UX 11i version 2 (at first release), Windows Server 2003 (at first release + 2 to 4 months) and Linux (first release + 6 months) operating systems.

Each npar has its own independent CPUs, memory and I/O resources consisting of the resources of the cells that make up the partition. Resources (cell boards and/or I/O chassis) may be removed from one npar and added to another without having to physically manipulate the hardware, but rather by using commands that are part of the System Management interface. The table below shows the maximum size of npars per operating system:

	HP UX 11i Version 2	Windows Server 2003	Red Hat Enterprise Linux AS 3 or Debian Linux
Maximum size of npar	64 CPUs, 512 GB RAM	64 CPUs, 512 GB RAM	8 CPUs, 64 GB RAM
Maximum number of npars	16	16	16

For information on type of I/O cards for networking and mass storage for each operating environment, please refer to the Technical Specifications section. For licensing information for each operating system, please refer to the Ordering Guide.

Superdome supports static partitions. Static partitions imply that any npar configuration change requires a reboot of the npar. In a future HP-UX and Windows release, dynamic npars will be supported. Dynamic npars imply that npar configuration changes do not require a reboot of the npar. Using the related capabilities of dynamic reconfiguration (i.e. on-line addition, on-line removal), new resources may be added to an npar and failed modules may be removed and replaced while the npar continues in operation. Adding new npars to Superdome system does not require a reboot of the system.

Single System Reliability/Availability Features

Superdome high availability offering is as follows:

NOTE: Online addition/replacement for cell boards is not currently supported and will be available in a future HP-UX release. Online addition/replacement of individual CPUs and memory DIMMs will never be supported.)



Configuration

- **CPU:** The features below nearly eliminate the down time associated with CPU cache errors (which are the majority of CPU errors). If a CPU is exhibiting excessive cache errors, HP-UX 11i version 2 will ONLINE activate to take its place. Furthermore, the CPU cache will automatically be repaired on reboot, eliminating the need for a service call.
 - Dynamic processor resilience w/ iCOD enhancement.
NOTE: Dynamic processor resilience and iCOD are not supported when running Windows Server 2003 or Linux in the partition.
 - CPU cache ECC protection and automatic de allocation
 - CPU bus parity protection
 - Redundant DC conversion
- **Memory:** The memory subsystem design is such that a single SDRAM chip does not contribute more than 1 bit to each ECC word. Therefore, the only way to get a multiple bit memory error from SDRAMs is if more than one SDRAM failed at the same time (rare event). The system is also resilient to any cosmic ray or alpha particle strike because these failure modes can only affect multiple bits in a single SDRAM. If a location in memory is "bad", the physical page is de allocated dynamically and is replaced with a new page without any OS or application interruption. In addition, a combination of hardware and software scrubbing is used for memory. The software scrubber reads/writes all memory locations periodically. However, it does not have access to "locked down" pages. Therefore, a hardware memory scrubber is provided for full coverage. Finally data is protected by providing address/control parity protection.
 - Memory DRAM fault tolerance, i.e. recovery of a single SDRAM failure
 - DIMM address / control parity protection
 - Dynamic memory resilience, i.e. page de allocation of bad memory pages during operation.
NOTE: Dynamic memory resilience is not supported when running Windows Server 2003 or Linux in the partition.
 - Hardware and software memory scrubbing
 - Redundant DC conversion
 - Cell COD.
NOTE: Cell COD is not supported when Windows Server 2003 or Linux is running in the partition.
- **I/O:** Partitions configured with dual path I/O can be configured to have no shared components between them, thus preventing I/O cards from creating faults on other I/O paths. I/O cards in hardware partitions (npars) are fully isolated from I/O cards in other hard partitions. It is not possible for an I/O failure to propagate across hard partitions. It is possible to dynamically repair and add I/O cards to an existing running partition.
 - Full single wire error detection and correction on I/O links
 - I/O cards fully isolated from each other
 - HW for the Prevention of silent corruption of data going to I/O
 - On line addition/replacement (OLAR) for individual I/O cards, some external peripherals, SUB/HUB.
 - Parity protected I/O paths
 - Dual path I/O
- **Crossbar and Cabinet Infrastructure:**
 - Recovery of a single crossbar wire failure
 - Localization of crossbar failures to the partitions using the link
 - Automatic de-allocation of bad crossbar link upon boot
 - Redundant and hotswap DC converters for the crossbar backplane
 - ASIC full burn-in and "high quality" production process
 - Full "test to failure" and accelerated life testing on all critical assemblies
 - Strong emphasis on quality for multiple-nPartition single points of failure (SPOFs)
 - System resilience to Management Processor (MP)
 - Isolation of nPartition failure
 - Protection of nPartitions against spurious interrupts or memory corruption
 - Hot swap redundant fans (main and I/O) and power supplies (main and backplane power bricks)
 - Dual power source
 - Phone-Home capability

HA Cluster-In-A-Box" Configuration: The "HA Cluster-In-A-Box" allows for failover of users' applications between hardware partitions (npars) on a single Superdome system. All providers of mission critical solutions agree that failover between clustered systems provides the safest availability-no single points of failures (SPOFs) and no ability to propagate failures between systems. However, HP supports the configuration of HA cluster software in a single system to allow the highest possible availability for those users that need the benefits of a non-clustered solution, such as scalability and manageability. Superdome with this configuration will provide the greatest single system availability configurable. Since no single system solution in the industry provides protection against a SPOF, users that still need this kind of safety and HP's highest availability should use HA cluster software in a multiple system HA configuration. Multiple HA software clusters can be configured within a single Superdome system (i.e., two 4 -ode clusters configured within a 32-way Superdome system).

Configuration

- HP-UX: Serviceguard and Serviceguard Extension for RAC
- Windows Server 2003: Microsoft Cluster Service (MSCS) - limited configurations supported
- Linux: Serviceguard for Linu

Multi-system High Availability

HP-UX 11i v2:

Any Superdome partition that is protected by Serviceguard or Serviceguard Extension for RAC can be configured in a cluster with:

- Another Superdome with Itanium 2 processors
- One or more standalone non Superdome systems with Itanium 2 processors
- Another partition within the same single cabinet Superdome (refer to "HA Cluster in a Box" above for specific requirements)

Separate partitions within the same Superdome system can be configured as part of different Serviceguard clusters.

Geographically Dispersed Cluster Configurations

HP-UX 11i v2:

The following Geographically Dispersed Cluster solutions fully support cluster configurations using Superdome systems. The existing configuration requirements for non Superdome systems also apply to configurations that include Superdome systems. An additional recommendation, when possible, is to configure the nodes of cluster in each datacenter within multiple cabinets to allow for local failover in the case of a single cabinet failure. Local failover is always preferred over a remote failover to the other datacenter. The importance of this recommendation increases as the geographic distance between datacenters increases.

- Extended Campus Clusters (using Serviceguard with MirrorDisk/UX)
- MetroCluster with Continuous Access XP
- MetroCluster with EMC SRDF
- ContinentalClusters

From an HA perspective, it is always better to have the nodes of an HA cluster spread across as many system cabinets (Superdome and non Superdome systems) as possible. This approach maximizes redundancy to further reduce the chance of a failure causing down time.

Windows Server 2003

Any Superdome partition that is protected by Microsoft Cluster Service for Windows Server 2003, Datacenter Edition can be configured in a cluster of up to 8 nodes with:

- Another Superdome complex
- Another partition within the same single cabinet Superdome with an identical hardware configuration

Furthermore, geographically dispersed clusters are supported utilizing a single quorum resource (Cluster Extension XP for Windows). Specific Superdome Windows Server 2003 cluster configurations will be announced later in 2003.

Linux

Support of Serviceguard on Linux and Cluster Extension on Linux should be available in late 2003.

Supportability Features

Superdome now supports the Console and Support Management Station in one device.

Console Access (Management Processor [MP])

The optimal configuration of console device(s) depends on a number of factors, including the customer's data center layout, console security needs, customer engineer access needs, and the degree with which an operator must interact with server or peripheral hardware and a partition (i.e. changing disks, tapes). This section provides a few guidelines. However the configuration that makes best sense should be designed as part of site preparation, after consulting with the customer's system administration staff and the field engineering staff.

Customer data centers exhibit a wide range of configurations in terms of the preferred physical location of the console device. (The term "console device" refers to the physical screen/keyboard/mouse that administrators and field engineers use to access and control the server.) The Superdome server enables many different configurations by its flexible configuration of

Configuration

access to the MP, and by its support for multiple geographically distributed console devices.

Three common data center styles are:

- The secure site where both the system and its console are physically secured in a small area.
- The "glass room" configuration where all the systems' consoles are clustered in a location physically near the machine room.
- The geographically dispersed site, where operators administer systems from consoles in remote offices.

These can each drive different solutions to the console access requirement.

The considerations listed below apply to the design of provision of console access to the server. These must be considered during site preparation.

- The Superdome server can be operated from a VT100 or an hpterm compatible terminal emulator. However some programs (including some of those used by field engineers) have a more friendly user interface when operated from an hpterm.
- LAN console device users connect to the MP (and thence to the console) using terminal emulators that establish telnet connections to the MP. The console device(s) can be anywhere on the network connected to either port of the MP.
- Telnet data is sent between the client console device and the MP "in the clear", i.e. unencrypted. This may be a concern for some customers, and may dictate special LAN configurations.
- If an HP-UX workstation is used as a console device, an hpterm window running telnet is the recommended way to connect to the MP. If a PC is used as a console device, Reflection1 configured for hpterm emulation and telnet connection is the recommended way to connect to the MP.
- The MP currently supports a maximum of 16 telnet-connected users at any one time.
- It is desirable, and sometimes essential for rapid time to repair to provide a reliable way to get console access that is physically close to the server, so that someone working on the server hardware can get immediate access to the results of their actions. There are a few options to achieve this:
 - Place a console device close to the server.
 - Ask the field engineer to carry in a laptop, or to walk to the operations center.

Use a system that is already in close proximity of the server such as the Instant Support Enterprise Edition (ISEE) or the System Management Station as a console device close to the system.

- The system administrator is likely to want to run X applications or a browser using the same client that they access the MP and partition consoles with. This is because the partition configuration tool, parmgr, has a graphical interface. The system administrator's console device(s) should have X window or browser capability, and should be connected to the system LAN of one or more partitions.

Functional capabilities:

- Local console physical connection (RS-232)
- Display of system status on the console (Front panel display messages)
- Console mirroring between LAN and RS-232 ports
- System hard and soft (TOC or INIT) reset capability from the console.
- Password secured access to the console functionality
- Support of generic terminals (i.e. VT100 compatible).
- Power supply control and monitoring from the console. It will be possible to get power supply status and to switch power on/off from the console.
- Console over the LAN. This means that a PC or HP workstation can become the system console if properly connected on the customer LAN. This feature becomes especially important because of the remote power management capability. The LAN will be implemented on a separate port, distinct from the system LAN, and provide TCP/IP and Telnet access.
- There is one MP per Superdome cabinet, thus there are two (2) for Superdome 64-way. But one, and only one, can be active at a time. There is no redundancy or failover feature.

Configuration

Windows

For Windows Server 2003 customers desiring full visibility to the Superdome Windows partition, an IP console solution is available to view the partition while the OS is rebooting (in addition to normal Windows desktop). Windows Terminal Services (standard in Windows Server 2003) can provide remote access, but does not display VGA during reboot. For customers who mandate VGA access during reboot, the IP console switch (262586-B21), used in conjunction with a VGA/USB card in the partition (A6869A) is the solution.

In order to have full graphical console access when running Windows Server 2003 on Superdome, the 3×1×16 IP Console Switch (product number 262586-B21) is required.

The features of this switch are as follows:

- Provides keyboard, video and mouse (KVM) connections to 16 direct attached Windows partitions (or servers) - expandable to 128.
- Allows access to partitions (or servers) from a remote centralized console.
- 1 for local KVM
- 3 concurrent remote users (secure SSL data transfer across network)
- Single screen switch management with the IP Console Viewer Software:
 - Authentication
 - Administration
 - Client Software

If the full graphical console access is needed, the following must be ordered.

Component	Product Number
3×1×16 IP console switch (100 240V)-1 switch per 16 OS instances (n≤16), each connected to VGA card	262586-B21
8 to 1 console expander-Order expander if there are more than 16 OS instances	262589-B21
USB interface adapters-Order one per OS instance	336057-001
CAT5 cable-Order one per OS instance	C7542A
AB243A 1U KVM-For local KVM - provides a full 15" digital display, keyboard, mouse and console switch in only 1U	221546-B21

Support Management Station

The purpose of the Support Management Station (SMS) is to provide Customer Engineers with an industry-leading set of support tools, and thereby enable faster troubleshooting and more precise problem root cause analysis. It also enables remote support by factory experts who consult with and back up the HP Customer Engineer. The SMS complements the proactive role of HP's Instant Support Enterprise Edition (ISEE) (which is offered to Mission Critical customers), by focusing on reactive diagnosis, for both mission critical and non mission critical Superdome customers.

The user of the SMS is the HP Customer Engineer and HP Factory Support Engineer. The Superdome customer benefits from their use of the SMS by receiving faster return to normal operation of their Superdome server, and improved accuracy of fault diagnosis, resulting in fewer callbacks. HP can offer better service through reduced installation time.

Only one SMS is required per customer site (or data center) connected to each platform via Ethernet LAN. Physically, it would be beneficial to have the SMS close to the associated platforms because the customer engineer will run the scan tools and would need to be near platform to replace failing hardware. The physical connection from the platform is an Ethernet connection and thus, the absolute maximum distance is not limited by physical constraints.

The SMS supports a single LAN interface that is connected to the Superdome and to the customer's management LAN. When connected in this manner, SMS operations can be performed remotely.

Physical Connection:

The SMS will contain one physical Ethernet connection, namely a 10/100Base-T connection. Note that the connection on Superdome (MP) is also 10/100Base-T, as is the LAN connection on the core I/O card installed in each hardware partition.

For connecting more than one Superdome server to the SMS, a LAN hub is required for the RJ-45 connection. A point to point connection is only required for one Superdome server to one SMS.

Configuration

Functional Capabilities:

- Allows local access to SMS by CE.
- Provides integrated console access, providing hpterm emulation over telnet and web browser, connecting over LAN or serial to a Superdome system
- Provides remote access over a LAN or dialup connection:
 - ftp server with capability to ftp the firmware files and logs
 - dialup modem access support (e.g. PC Anywhere or VNC)
- Provides seamless integration with data center level management.
- Provides partition logon capability, providing hpterm emulation over telnet, X-windows, and Windows Terminal Services capabilities.
- Provides following diagnostics tools:
 - Runs HP's proven highly effective JTAG scan diagnostic tools, which offer rapid fault resolution to the failing wire.
 - Superdome HPMC and MCA analyzer
 - Console log storage and viewing
 - Event log storage and viewing
 - Partition and memory adviser flash applications
- Supports updating platform and system firmware.
- Always on event and console logging for Superdome systems, which captures and stores very long event and console histories, and allows HP specialists to analyze the first occurrence of a problem.
- Allows more than one LAN connected response center engineer to look at SMS logs simultaneously.
- Can be disconnected from the Superdome systems and not disrupt their operation.
- Provides ability to connect a new Superdome system to the SMS and be recognized by scan software.
- Scans one Superdome system while other Superdome systems are connected (and not disrupt the operational systems).
- Supports multiple, heterogeneous Superdome platforms.

Minimum Hardware Requirements:

The SMS should meet the following **minimum** hardware requirements:

- ProLiant ML350 G3 running Windows 2000 Server SP3 including:
 - Modem
 - DVD R/W
 - Keyboard/monitor/mouse
 - 512-MB memory
 - Options:
 - Factory racked or field racked
 - Rack mount or desk mount keyboard/monitor/mouse/platform (bundled CPL line items)
- NOTE:** The rack-mount option of the SMS will not be available for ordering until July 1, 2003.

Software Requirements:

The SMS will run Windows 2000 SP3 as the default operating system. The SMS will follow the Windows OS roadmap and support later versions of this operating system as needed.

NOTE: The CE Tool is used by the CE to service the system and is not part of the purchased system.

System Management Features

HP-UX

- **HP-UX Servicecontrol Manager** is the central point of administration for management applications that address the configuration, fault, and workload management requirements of an adaptive infrastructure.
- **Servicecontrol Manager** maintains both effective and efficient management of computing resources. It integrates with many other HP-UX specific system management tools, including the following, which are available on Itanium 2 based servers:
 - **Ignite-UX**-Ignite-UX addresses the need for HP-UX system administrators to perform fast deployment for one or many servers. It provides the means for creating and reusing standard system configurations, enables replication of systems, permits post installation customizations, and is capable of both interactive and unattended operating modes.
 - **Software Distributor (SD)** is the HP-UX administration tool set used to deliver and maintain HP-UX operating systems and layered software applications. Delivered as part of HP-UX, SD can help you manage your HP-UX operating system, patches, and application software on HP Itanium2 based servers.
- **System Administration Manager (SAM)** is used to manage accounts for users and groups, perform auditing

Configuration

and security, and handle disk and file system management and peripheral device management. Servicecontrol Manager enables these tasks to be distributed to multiple systems and delegated using role based security.

HP-UX Kernel Configuration-for self optimizing kernel changes. The new HP-UX Kernel Configuration tool allows users to tune both dynamic and static kernel parameters quickly and easily from a Web based GUI to optimize system performance. This tool also sets kernel parameter alarms that notify you when system usage levels exceed thresholds.

- **Partition Manager** creates and manages nPartitions-hard partitions for high end servers. Once the partitions are created, the systems running on those partitions can be managed consistently with all the other tools integrated into Servicecontrol Manager. Key features include:
 - Easy-to-use, familiar graphical user interface.
 - Runs locally on a partition, or remotely. The Partition Manager application can be run remotely on any system running HP-UX 11i Version 2 and eventually select Windows releases and remotely manage a complex either by 1) communicating with a booted OS on an nPartition in the target complex via WBEM, or 2) communicating with the service processor in the target complex via IPMI over LAN. The latter is especially significant because a complex can be managed with NONE of the nPartitions booted.
 - Full support for creating, modifying, and deleting hardware partitions.
 - Automatic detection of configuration and hardware problems.
 - Ability to view and print hardware inventory and status.
 - Big picture views that allow system administrators to graphically view the resources in a server and the partitions that the resources are assigned to.
 - Complete interface for the addition and replacement of PCI devices.
 - Comprehensive online help system.
- **Security Patch Check** determines how current a systems security patches are, recommends patches for continuing security vulnerabilities and warns administrators about recalled patches still present on the system.
- **System Inventory Manager** is for change and asset management. It allows you to easily collect, store and manage inventory and configuration information for HP-UX based servers. It provides an easy to use, Web based interface, superior performance, and comprehensive reporting capabilities
- **Event Monitoring Service (EMS)** keeps the administrator of multiple systems aware of system operation throughout the cluster, and notifies the administrator of potential hardware or software problems before they occur. HP Servicecontrol Manager can launch the EMS interface and configure EMS monitors for any node or node group that belongs to the cluster, resulting in increased reliability and reduced downtime.
- **Process Resource Manager (PRM)** controls the resources that processes use during peak system load. PRM can manage the allocation of CPU, memory resources, and disk bandwidth. It allows administrators to run multiple mission critical applications on a single system, improve response time for critical users and applications, allocate resources on shared servers based on departmental budget contributions, provide applications with total resource isolation, and dynamically change configuration at any time-even under load. (fee based)
- **HP-UX Workload Manager (WLM)** A key differentiator in the HP-UX family of management tools, Workload Manager provides automatic CPU resource allocation and application performance management based on prioritized service level objectives (SLOs). In addition, WLM allows administrators to set real memory and disk bandwidth entitlements (guaranteed minimums) to fixed levels in the configuration. The use of workload groups and SLOs improves response time for critical users, allows system consolidation, and helps manage user expectations for performance. (Fee based)
- **HP's Management Processor** enables remote server management over the Web regardless of the system state. In the unlikely event that none of the nPartitions are booted, the Management Processor can be accessed to power cycle the server, view event logs and status logs, enable console redirection, and more. The Management Processor is embedded into the server and does not take a PCI slot. And, because secure access to the Management Processor is available through SSL encryption, customers can be confident that its powerful capabilities will be available only to authorized administrators. New features that will be available include:
 - Support for Web Console that provides secure text mode access to the management processor
 - Reporting of error events from system firmware.
 - Ability to trigger the task of PCI OL* from the management processor.
 - Ability to scan a cell board while the system is running. (only available for partitionable systems)
 - Implementation of management processor commands for security across partitions so that partitions do not modify system configuration (only available for partitionable systems).
- **OpenView Operations Agent**-collects and correlates OS and application events (fee based)
- **OpenView Performance Agent**-determines OS and application performance trends (fee based)
- **OpenView GlancePlus**-shows real time OS and application availability and performance data to diagnose problems (fee based)
- **OpenView Data Protector (Omniback II)**-backs up and recovers data (fee based)

In addition, the Network Node Manager (NNM) management station will run on HP-UX Itanium 2 based servers. NNM automatically discovers, draws (maps), and monitors networks and the systems connected to them.

Configuration

All other OpenView management tools, such as OpenView Operations, Service Desk, and Service Reporter, will be able to collect and process information from the agents running on Itanium 2-based servers running HP-UX.

Windows Server 2003, Datacenter Edition

- **The HP Essentials Foundation Pack for Windows** is a complete toolset to install, configure, and manage Itanium2 servers running Windows. Included in the Pack is the Smart Setup DVD which contains all the latest tested and compatible HP Windows drivers, HP firmware, HP Windows utilities, and HP management agents that assist in the server deployment process by preparing the server for installation of standard Windows operating system and in the on going management of the server. Please note that this is available for HP service personnel but not provided to end customers.
- **Partition Manager** creates and manages nPartitions-hard partitions for high end servers. Once the hard partitions are created, the Windows Server 2003 resources running on those partitions can be managed consistently with the Windows System Resource Manager and Insight Manager through the System Management Homepage (see below). Key features include full support for creating, modifying, and deleting hardware partitions.
NOTE: At first release, Partition Manager will require a PC SMS running Partition Manager Command Line (A9801A or A9802A) or an HP-UX 11i Version 2 partition or separate device (i.e. Itanium2-based workstation or server running HP-UX 11i Version 2) in order to configure Windows partitions. Refer to HP-UX section above for key features of Partition Manager.
- **Insight Manager 7** maximizes system uptime and provides powerful monitoring and control. Insight Manager 7 delivers pre failure alerting for servers ensuring potential server failures are detected before they result in unplanned system downtime. Insight Manager 7 also provides inventory reporting capabilities that dramatically reduce the time and effort required to track server assets and helps systems administrators make educated decisions about which system may require hardware upgrades or replacement. And Insight Manager 7 is an effective tool for managing your HP desktops and notebooks as well as non HP devices instrumented to SNMP or DMI.
- **System Management Homepage** displays critical management information through a simple, task oriented user interface. All system faults and major subsystem status are now reported within the initial System Management Homepage view. In addition, the new tab-based interface and menu structure provide one click access to server log. The System Management Homepage is accessible either directly through a browser (with the partition's IP address) or through a management application such as Insight Manager 7 or an enterprise management application.
- **HP's Management Processor** enables remote server management over the Web regardless of the system state. In the unlikely event that the operating system is not running, the Management Processor can be accessed to power cycle the server, view event logs and status logs, enable console redirection, and more. The Management Processor is embedded into the server and does not take a PCI slot. And, because secure access to the Management Processor is available through SSL encryption, customers can be confident that its powerful capabilities will be available only to authorized administrators. New features on the management processor include:
 - Support for Web Console that provides secure text mode access to the management processor
 - Reporting of error events from system firmware.
 - Ability to trigger the task of PCI OL* from the management processor.
 - Ability to scan a cell board while the system is running.
 - Implementation of management processor commands for security across partitions so that partitions do not modify system configuration.
- **OpenView Management Tools**, such as OpenView Operations and Network Node Manager, will be able to collect and process information from the SNMP agents and WMI running on Windows Itanium 2 based servers. In the future, OpenView agents will be able to directly collect and correlate event, storage, and performance data from Windows Itanium 2 based servers, thus enhancing the information OpenView management tools will process and present.

Linux

- **Insight Manager 7** maximizes system uptime and provides powerful monitoring and control. Insight Manager 7 also provides inventory reporting capabilities that dramatically reduce the time and effort required to track server assets and helps systems administrators make educated decisions about which system may require hardware upgrades or replacement. And Insight Manager 7 is an effective tool for managing your HP desktops and notebooks as well as non HP devices instrumented to SNMP or DMI.
- **The HP Enablement Kit for Linux** facilitates setup and configuration of the operating system. This kit includes System Imager, an open source operating system deployment tool. System Imager is a golden image based tool and can be used for initial deployment as well as updates.
Partition Manager creates and manages nPartitions-hard partitions for high-end servers. Once the partitions are created, the systems running on those partitions can be managed consistently with all the other tools integrated into Servicecontrol Manager.
NOTE: At first release, Partition Manager will require an HP-UX 11i Version 2 partition or separate device (i.e.

Configuration

Itanium2 based workstation or server running HP-UX 11i Version 2) in order to configure Linux partitions. Refer to HP-UX section above for key features of Partition Manager.

- **HP's Management Processor** enables remote server management over the Web regardless of the system state. In the unlikely event that the operating system is not running, the Management Processor can be accessed to power cycle the server, view event logs and status logs, enable console redirection, and more. The Management Processor is embedded into the server and does not take a PCI slot. And, because secure access to the Management Processor is available through SSL encryption, customers can be confident that its powerful capabilities will be available only to authorized administrators.
 - Support for Web Console that provides secure text mode access to the management processor
 - Reporting of error events from system firmware.
 - Ability to trigger the task of PCI OL* from the management processor.
 - Ability to scan a cell board while the system is running. (only available for partitionable systems)
 - Implementation of management processor commands for security across partitions so that partitions do not modify system configuration. (only available for partitionable systems)

General Site Preparation AC Power Requirements Rules

The modular, N+1 power shelf assembly is called the Front End Power Subsystem (FEPS). The redundancy of the FEPS is achieved with 6 internal Bulk Power Supplies (BPS), any five of which can support the load and performance requirements.

Input Options

Reference the Site Preparation Guide for detailed power configuration options.

Input Power Options

PDCA Product Number	Source Type	Source Voltage (nominal)	PDCA Required	Input Current Per Phase 200-240 VAC	Power Required
A5800A Option 006	3-phase	Voltage range 200- 240 VAC, phase-to- phase, 50/60 Hz	4-wire	44 A Maximum per phase	2.5 meter UL power cord and OL approved plug provided. The customer must provide the mating in line connector or purchase quantity one A6440A opt 401 to receive a mating in line connector. An electrician must hardwire the in- line connector to 60 A/63 A site power. ^{a,b,c}
A5800A Option 007	3-phase	Voltage range 200- 240 VAC, phase-to- neutral, 50/60 Hz	5-wire	24 A Maximum per phase	2.5 meter <HAR> power cord and VDE approved plug provided. The customer must provide the mating in line connector or purchase quantity 1 A6440A opt 501 to receive a mating in line connector. An electrician must hardwire the in-line connector to 30 A/32 A site power. ^{a,b,d}

a. A dedicated branch is required for each PDCA installed.

b. In the U.S.A, site power is 60 Amps; in Europe site power is 63 Amps.

c. Refer to the Option 006 and 007 Specifics Table for detailed specifics related to this option.

d. In the U.S.A. site power is 30 Amps; in Europe site power is 32 Amps.

Option 006 and 007 Specifics^a

Configuration

PDCA Product Number	Attached Power Cord	Attached Plug	Customer Provided Part	
			In Line Connector	Panel Mount Receptacle
A5800A Option 006	OLFLEX 190 (PN 600804), four-conductor, 6-AWG (16 mm ²), 600-Volt, 60-Amp, 90-degree C, UL and CSA approved, conforms to CE directives GN/YW ground wire.	Mennekes ME 460P9 3-phase, 4-wire, 60-Amp, 250-Volt, UL-approved. Color blue, IEC 309-1, IEC 309-1, grounded at 3:00 o'clock.	Mennekes ME 460C9 3-phase, 4-wire, 60-amp, 250-Volt, UL-approved. Color blue, IEC 309-1, IEC 309-1, grounded at 9:00 o'clock. ^a	Mennekes ME 460R9 3-phase, 4-wire, 60-amp, 250-Volt, UL-approved. Color blue, IEC 309-1, IEC 309-1, grounded at 9:00 o'clock. ^b
A5800A Option 007	Five conductors, 10-AWG (6 mm ²), 450/475-Volt, 32-Amp, <HAR> European wire cordage, GN/YW ground wire.	Mennekes ME 532P6-14 3-phase, 5-wire, 32-Amp, 450/475-volt, VDE-certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00 o'clock.	Mennekes ME 532C6 16 3-phase, 5-wire, 32-Amp, 450/475-Volt, VDE-certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00 o'clock. ^c	Mennekes ME532R6 1276 3-phase, 5-wire, 32-Amp, 450/475-Volt, VDE-certified, color red, IEC 309-1, IEC 309-2, grounded at 6:00 o'clock. ^b

a. In line connector is available from HP by purchasing A6440A, Option 401.

b. Panel mount receptacles must be purchased by the customer from a local Mennekes supplier.

c. In line connector is available from HP by purchasing A6440A, Option 501.

NOTE: A qualified electrician must wire the PDCA in line connector to site power using copper wire and in compliance with all local codes.

Input Requirements

Reference the Site Preparation Guide for detailed power configuration requirements.

Requirements	Value	Conditions/Comments
Nominal Input Voltage (VAC rms)	200/208/220/230/240	
Input Voltage Range (VAC rms)	200-240	Auto selecting. Measure at input terminals
Frequency Range (Hz)	50/60	
Number of Phases	3	3-phase 5-wire with power cord; 3-phase 4-wire with power cord
Maximum Input Current (A rms), 3-Phase 5-wire	20	3-phase source with a source voltage of 220 VAC measured phase to neutral
Maximum Input Current (A rms), 3-Phase 4-wire	40	3-phase source with a source voltage of either 208 VAC or 230 VAC measured phase to phase
Maximum Inrush Current (A peak)	90	
Circuit Breaker Rating (A), 3-Phase 5-wire	25 A	Per phase
Circuit Breaker Rating (A), 3-Phase 4-wire	45 A	Per phase
Power Factor Correction	0.95 minimum	
Ground Leakage Current (mA)	>3.5 mA, with 6 BPSs installed	Warning label applied to the PDCA at the AC Mains input

Cooling Requirements

- The cooling system in Superdome was designed to maintain reliable operation of the system in the specified environment. In addition, the system is designed to provide redundant cooling (i.e. N+1 fans and blowers) that allows all of the cooling components to be "hot swapped."
Superdome was designed to operate in all data center environments with any traditional room cooling scheme (i.e. raised floor environments) but in some cases where data centers have previously installed high power density systems,

Configuration

alternative cooling solutions may need to be explored by the customer. HP has teamed with Liebert to develop an innovative data room cooling solution called DataCool. DataCool is a patented overhead climate system utilizing fluid based cooling coils and localized blowers capable of cooling heat loads of several hundred watts per square foot. Some of DataCool's highlights are listed below:

- Liebert has filed for several patents on DataCool
- DataCool, based on Liebert's TeleCool, is an innovative approach to data room cooling
- Liquid cooling heat exchangers provide distributed cooling at the point of use
- Delivers even cooling throughout the data center preventing hot spots
- Capable of high heat removal rates (500 W per square foot)
- Floor space occupied by traditional cooling systems becomes available for revenue generating equipment.
- Enables cooling upgrades when installed in data rooms equipped with raised floor cooling

DataCool is a custom engineered overhead solution for both new data center construction and for data room upgrades for high heat loads. It is based on Liebert's TeleCool product, which has been installed in 600 telecommunications equipment rooms throughout the world. The system utilizes heat exchanger pump units to distribute fluid in a closed system through patented cooling coils throughout the data center. The overhead cooling coils are highly efficient heat exchangers with blowers that direct the cooling where it is needed. The blowers are adjustable to allow flexibility for changing equipment placement or room configurations. Equipment is protected from possible leaks in the cooling coils by the patented monitoring system and purge function that detects any leak and safely purges all fluid from the affected coils. DataCool has interleaved cooling coils to enable the system to withstand a single point of failure and maintain cooling capability.

Features and Benefits

- Fully distributed cooling with localized distribution
- Even cooling over long distances
- High heat load cooling capacity (up to 500 W per square foot)
- Meets demand for narrow operating temperature for computing systems
- Allows computer equipment upgrade for existing floor cooled data rooms
- Floor space savings from removal of centralized air distribution
- Withstand single point of failures

For More Information

http://www.liebert.com/assets/products/english/products/env/datacool/60hz/bro_8pg/acrobat/sl_16700.pdf

HP has entered into an agreement with Liebert to reference sell the DataCool solution.

- The HP/Liebert business relationship is managed by the HP Complementary Products Division.
- DataCool will be reference by HP. Liebert will perform installation, service and support.
- HP will compensate the HP Sales Representative and District Manager for each DataCool that Liebert sells to a customer referred by HP.
- An HP/Liebert DataCool website will be setup to get more information on the product and to manage the reference sales process. Please go to <http://hpcp.grenoble.hp.com/> for more information.

Environmental

- 68 to 86 degrees F (20 to 30 degrees C) inlet ambient temperature
- 0 to 10,000 feet (0 to 3048 meters)
- 2600 CFM with N+1 blowers. 2250 CFM with N.
- 65 dBA noise level

Uninterruptible Power Supplies (UPS)

HP will be reselling high-end (10-kW and above) three-phase UPS systems from our partners. We will test and qualify a three-phase UPS for Superdome. The UPS is planned to be available Q1 FY01.

- All third-party UPS resold by HP will be tested and qualified by HP to ensure interoperability with our systems
- We plan to include ups_mond ups communications capability in the third party UPS(s), thus ensuring consistent communications strategy with our PowerTrust UPS(s)
- We will also establish a support strategy with our third-party UPS partners to ensure the appropriate level of support our customer have come to expect from HP.

Configuration

- For more information on the product and to manage the reference sales process. Please go to <http://hpcp.grenoble.hp.com/> for more information.

APC Uninterruptible Power Supplies for Superdome

The Superdome team has qualified the APC Silcon 3-phase 20 kW UPS for Superdome.

There are several configurations that can be utilized depending on the Superdome configuration your customer is deploying. They range from a 64-way Superdome with dual cord and dual UPS with main tie main to a 32-way Superdome with single cord and single UPS. In all configurations the APC Silcon SL20KFB2 has been tested and qualified by the Superdome engineers to ensure interoperability.

HP UPS Solutions

Product Number/Description	Quantity/Configuration	Watt	VA	Technology	Family	Package	Output
SL20KFB2 APC Silcon 3-phase UPS	<ul style="list-style-type: none"> Quantity 2/ 32- or 64-way dual-cord/dual-UPS with main-tie-main Quantity 1/ 32- or 64-way single-cord/ single-UPS 	20 kW	20 kVA	Delta conversion on line double conversion	APC Silcon 3-phase	Standalone rack	Configurable for 200: 208 or 220V 3 phase nominal output voltage
QJB22830 Switch Gear	<ul style="list-style-type: none"> Quantity 1/ 32- or 64-way dual- cord/dual-UPS with main-tie-main Quantity 0/ 32- or 64-way single- cord/ single-UPS 	N/A	N/A	N/A	Customer Design for Superdome	N/A	N/A
WSTRUP5X8- SL10 Start Up Service	<ul style="list-style-type: none"> Quantity 2/ 32- or 64-way dual-cord/dual-UPS with main-tie-main Quantity 1/ 32- or 64-way single-cord/ single-UPS 	N/A	N/A	N/A	N/A	N/A	N/A
WONSITENBD- SL10 Next Business Day On site Service	<ul style="list-style-type: none"> Quantity 2/ 32- or 64-way dual- cord/dual-UPS with main-tie-main Quantity 1/ 32- or 64-way single- cord/ single-UPS 	N/A	N/A	N/A	N/A	N/A	N/A

Power Protection

Runtimes

The UPS will provide battery backup to allow for a graceful shutdown in the event of a power failure. Typical runtime on the APC SL20KFB2 Silcon 3 Phase UPS varies with the kW rating and the load. The APC SL20KFB2 UPS provides a typical runtime of 36.7 minutes at half load and 10.7 at full load. If additional run time is needed please contact your APC representative

Power Conditioning

Configuration

The APC SL20KFB2 provides unparalleled power conditioning with its Delta-Conversion on-line double conversion technology. This is especially helpful in regions where power is unstable.

Continuous Power during Short Interruptions of Input Power

The APC SL20KFB2 will provide battery backup to allow for continuous power to the connected equipment in the event of a brief interruption in the input power to the UPS. Transaction activity will continue during brief power outage periods as long as qualified UPS units are used to provide backup power to the SPU, the Expansion Modules, and all disk and disk array products.

UPS Configuration Guidelines

In general, the sum of the "Watt rating for UPS sizing" for all of the connected equipment should not exceed the watt rating of the UPS from which they all draw power. In previous configuration guides, this variable was called the "VA rating for UPS sizing." With Unity Power Factor, the Watt rating was the same as the kVA rating, so it didn't matter which one we used. VA is calculated by multiplying the voltage times the current. Watts, which is a measurement of true power, may be less than VA if the current and voltage are not in phase. APC SL20KFB2 has Unity Power Factor correction, so the kW rating equals the kVA rating. Be sure to add in the needs for the other peripherals and connected equipment. When sizing the UPS, allow for future growth as well. If the configuration guide or data sheet of the equipment you want to protect gives a VA rating, use this as the watt rating. If the UPS does not provide enough power for the additional devices such as system console and mass storage devices, additional UPSs may be required.

Superdome

The only qualified UPS available for use with Superdome is the APC SL20KFB2 Silcon 3 Phase 20-kW UPS.

The APC SL20KFB2 can provide power protection for the SPU and peripherals. If the system console and primary mass storage devices also require power protection (which is highly recommended) they may require one or more additional UPSs depending on the total Watts. Make sure that the total watts do not exceed the UPS's voltage rating.

Integration/Installation

The APC SL20KFB2 includes both field integration start up service and next day on site service for one year provided by APC.

Power Connections with the APC SL20KFB2

Product Number	Watts	NOM Out	Output Receptacles	Input Receptacles
SL20KFB2	20 kW	115/200 3PH, 120/208 3PH, 127/220 3PHV	Hardwire	Hardwire

Communications Connections

A DB-25 RS-232 Contact Closure connection is standard on all APC SL20KFB2 UPS. A WEB/SNMP card is also included.

Power Management

Description	Network interface cards that provide standards-based remote management of UPSs
General Features	Boot-P support, Built-in Web/SNMP management, Event logging, Flash Upgradeable, MD5 Authentication Security, Password Security, SNMP Management, Telnet Management, Web Management
Includes	CD with software, User Manual
Documentation	User Manual Installation Guide

Type of UPSs

Some customers may experience chronic "brown-out" situations or have power sources that are consistently at the lower spectrum of the standard voltage range. For example, the AC power may come in consistently at 92 VAC in a 110 VAC area. Heavy-load electrical equipment or power rationing are some of the reasons these situations arise. The APC SL20KFB2 units are designed to kick in before the AC power drops below the operating range of the HP Superdome Enterprise Server. Therefore, these UPS units may run on battery frequently if the AC power source consistently dips below the threshold voltage. This may result in frequent

system shutdowns and will eventually wear out the battery. Although the on-line units can compensate for the AC power shortfall, the battery life may be shortened. The best solution is to use a good quality boost transformer to "correct" the power source before it enters the UPS unit.

Configuration

Ordering Guidelines

- The APC SL20KFB2 Silcon 3-phase UPS units may be ordered as part of a new Superdome system order or as a field upgrade to an existing system.
- For new systems order please contact Ron Seredian at APC by e-mail at rseredia@apcc.com during the Superdome pre-consulting phase. APC will coordinate with HP to ensure the UPS is installed to meet the Superdome installation schedule.
- For field upgrades please contact Ron Seredian at APC by e-mail at rseredia@apcc.com when you determine a customer is in need and/or interested in power protection for Superdome. APC will coordinate with the customer to ensure the UPS is installed to meet their requirements.
- Numerous options can be ordered to compliment APC SL20KFB2 Silcon 3-phase UPS units. Your APC consultant can review these option with you are you can visit the APC website at www.apcc.com

Power Redundancy

Superdome servers, by default, provide an additional power supply for N+1 protection. As a result, Superdome servers will continue to operate in the event of a single power supply failure. The failed power supply can be replaced without taking the system down.

Multi-cabinet Configurations

When configuring Superdome systems that consist of more then one cabinet and include I/O expansion cabinets, certain guidelines must be followed, specifically the I/O interface cabling between the Superdome cabinet and the I/O expansion cabinet can only cross one additional cabinet due to cable length restrictions.

Configuration Guidelines/Rules

Superdome Configuration Guidelines/Rules

Category	Rule Index	Rule Description
General	1	Every Superdome complex requires connectivity to a Support Management Station (SMS). The PC-based SMS also serves as the system console.
	2	Every cell in a Superdome complex must be assigned to a valid physical location.
CPU	3	All CPUs in a cell are the same type, same Front Side Bus (FSB) frequency, and same core frequency.
Memory	4	Configurations with 8, 16 and 32 DIMM slots are recommended (i.e. are fully qualified and offer the best bandwidth performance).
	5	Configurations with 4 and 24 DIMM slots are supported (i.e. are fully qualified, but don't necessarily offer the best bandwidth performance).
	6	DIMMs can be deallocated in 2 DIMM increments (to support HA).
	7	Mixed DIMM sizes within a cell board are supported, but only in separate Mbat interleaving groups.
	8	System orders from the factory provide mixed DIMM sizes in recommended configurations only.
	9	For system orders from the factory, the same memory configuration must be used for all cells within a partition.
	10	DIMMs in the same rank must have SDRAMs with the same number of banks and row and column bits.
	11	Size of memory within an interleave group must be power of 2.
	12	DIMMs within the same interleave group must be same size and have same number of banks, row bits, and column bits.
	13	There are currently no restrictions on mixing DIMMs (of the same type) with different vendor SDRAMs.
I/O	14	One cell in every partition must be connected to an I/O chassis that contains a Core I/O card, a card connected to boot media, a card connected to removable media, and a network card with a connected network.
	15	A partition cannot have more I/O chassis than it has active cells.

Configuration		
	16	Removable media device controller should be in slot 8 of the I/O chassis.
	17	Core I/O card must be in slot 0 of the I/O chassis.
	18	Boot device controller should be in slot 1 of the I/O chassis
	19	PCI-X high bandwidth I/O cards should be in the high bandwidth slots in the I/O chassis
	20	Every I/O card in an I/O chassis must be assigned to a valid physical location.
	21	Every I/O chassis in a Superdome complex must be assigned to a valid physical location
Performance	22	The amount of memory on a cell should be evenly divisible by 4 GB if using 512-MB DIMMs or 8 GB if using 1-GB DIMMs, i.e. 8, 16 or 32 DIMMs. The cell has four memory subsystems and each subsystem should have an echelon (2 DIMMs) populated. The loading order of the DIMMs alternates among the four subsystems. This rule provides maximum memory bandwidth on the cell, by equally populating all four memory subsystems.
	23	All cells in a partition should have the same number of processors.
	24	The number of active CPUs per cell should be balanced across the partition, however minor differences are OK. (Example: 4 active CPUs on one cell and three active CPUs on the second cell)
	25	If memory is going to be configured as fully interleaved, all cells in a partition should have the same amount of memory (symmetric memory loading). Asymmetrically distributed memory affects the interleaving of cache lines across the cells. Asymmetrically distributed memory can create memory regions that are non optimally interleaved. Applications whose memory pages land in memory interleaved across just one cell can see up to 16 times less bandwidth than ones whose pages are interleaved across all cells.
	26	If a partition contains 4 or fewer cells, all the cells should be linked to the same crossbar (quad) in order to eliminate bottlenecks and the sharing of crossbar bandwidth with other partitions. In each Superdome cabinet, slots 0, 1, 2 and 3 link to the same crossbar and slots 4, 5, 6 and 7 link to the same crossbar.
	27	A Core I/O card should not be selected as the main network interface to a partition. A Core I/O card is a PCI X 1X card that possibly produces lower performance than a comparable PCI X 2X card.
	28	The number of cells in a partition should be a power of two, i.e., 2, 4, 8, or 16. Optimal interleaving of memory across cells requires that the number of cells be a power of two. Building a partition that does not meet this requirement can create memory regions that are non optimally interleaved. Applications whose memory pages land in the memory that is interleaved across just one cell can experience up to 16 times less bandwidth than pages which are interleaved across all 16 cells.
	29	<p>Before consolidating partitions in a Superdome 32-way or 64-way system, the following link load calculation should be performed for each link between crossbars in the proposed partition.</p> <p>Links loads less then 1 are best. As the link load begins to approach 2 performance bottlenecks may occur.</p> <p>For crossbars X and Y $\text{Link Load} = Q_x * Q_y / Q_t / L$, where</p> <ul style="list-style-type: none"> - Q_x is the number of cells connected to crossbar X (quad) - Q_y is the number of cells connected to crossbar Y (quad) - Q_t is the total number of cells in the partition - L is the number of links between crossbar X and Y (2 for Superdome 32-way systems and 1 for Superdome 64-way systems)
	30	Maximum performance for optimal configurations (power of two cells, uniform memory across cells, power of two DIMM ranks per cell)
	31	(If rule #30 cannot be met, rule #31 is recommended) Non-power of two cells, but still uniform memory across cells, power of two DIMM ranks per cell, uniform type of DIMM.
	32	(If rule #30 or #31 cannot be met, rule #32 is recommended) Same amount of memory in each cell, but possibly different memory types in each cell (for instance, a two cell configuration with 8 512MB DIMMs in one cell, and 4 1GB DIMMs in the other). Differences in memory across different cells within the same partition should be minimal for the best performance.

Configuration

Configuration	33	Same amount of memory in each cell, but non optimal and/or mixed loading within a cell (for instance, a two cell configuration with 16 512MB DIMMs and 8 1GB DIMMs in each cell).
	34	Non-uniform amount of memory across cells (this needs to boot and run, but performance is whatever you get).
	35	For the same amount of total memory, best performance is with a larger number of smaller size DIMMs.
Single System High Availability	36	Each cell should have at least two active CPUs.
	37	Each cell should have at least 4 GB (8 DIMMs) of memory using 512-MB DIMMs and at least 8 GB of memory using 1-GB DIMMs.
	38	I/O chassis ownership must be localized as much as possible. One way is to assign I/O chassis to partitions in sequential order starting from INSIDE the single cabinet, then out to the I/O expansion cabinet 'owned' by the single cabinet.
	39	I/O expansion cabinets can be used only when the main system cabinet holds maximum number of I/O card cages. Thus, the cabinet must first be filled with I/O card cages before using an I/O expansion cabinet.
	40	Single cabinets connected to form a dual cabinet (using flex cables) should use a single I/O expansion cabinet if possible.
	41	Spread enough connections across as many I/O chassis as it takes to become 'redundant' in I/O chassis'. In other words, if an I/O chassis fails, the remaining chassis have enough connections to keep the system up and running, or in the worst case, have the ability to reboot with the connections to peripherals and networking intact.
	42	All SCSI cards are configured in the factory as unterminated. Any auto termination is defeated. If auto termination is not defeatable by hardware, the card is not used at first release. Terminated cable would be used for connection to the first external device. In the factory and for shipment, no cables are connected to the SCSI cards. In place of the terminated cable, a terminator is placed on the cable port to provide termination until the cable is attached. This is needed to allow HP-UX to boot. The customer does not need to order the terminators for these factory integrated SCSI cards, since the customer will probably discard them. The terminators are provided in the factory by use of constraint net logic.
	43	Partitions whose I/O chassis are contained within a single cabinet have higher availability than those partitions that have their I/O chassis spread across cabinets.
	44	A partition's core I/O chassis should go in a system cabinet, not an I/O expansion cabinet
	45	A partition should be connected to at least two I/O chassis containing Core I/O cards. This implies that all partitions should be at least 2 cells in size. The lowest number cell or I/O chassis is the 'root' cell; the second lowest number cell or I/O chassis combo in the partition is the 'backup root' cell.
	46	A partition should consist of at least two cells.
Multi System High Availability (Please also refer to Multi System High Availability section following this table)	47	Not more than one partition should span a cabinet or a crossbar link. When crossbar links are shared, the partition is more at risk relative to a crossbar failure that may bring down all the cells connected to it.
	48	Multi initiator support is required for Serviceguard. The A5149A adapter will be required prior to Ultra160 SCSI adapters (A6828A and A6829A) support of multi initiator environments (available February 2003).
Traditional Multi System High Availability	49	To configure a cluster with no SPOF, the membership must extend beyond a single cabinet. The cluster must be configured such that the failure of a single cabinet does not result in the failure of a majority of the nodes in the cluster. The cluster lock device must be powered independently of the cabinets containing the cluster nodes. Alternative cluster lock solution is the Quorum Service, which resides outside the Serviceguard cluster providing arbitration services.
	50	A cluster lock is required if the cluster is wholly contained within two single cabinets (i.e., two Superdome/ 16-way or 32-way systems or two Superdome/PA 8800 32-way or 64-way systems) or two dual cabinets (i.e. two Superdome/ 64-way systems or two Superdome/PA

Configuration		8800 128 way systems). This requirement is due to a possible 50% cluster failure.
	51	Serviceguard only supports cluster lock up to four nodes. Thus a two cabinet configuration is limited to four nodes (i.e., two nodes in one dual cabinet Superdome/ 64-way system or Superdome/PA 8800 128-way system and two nodes in another dual cabinet Superdome/ 64-way system or Superdome/PA 8800 128 way system). The Quorum Service can support up to 50 clusters or 100 nodes (can be arbitrator to both HP UX and Linux clusters).
	52	Two-cabinet configurations must evenly divide nodes between the cabinets (i.e. 3 and 1 is not a legal 4-node configuration).
	53	Cluster lock must be powered independently of either cabinet.
	54	Root volume mirrors must be on separate power circuits.
	55	Redundant heartbeat paths are required and can be accomplished by using either multiple heartbeat subnets or via standby interface cards.
	56	Redundant heartbeat paths should be configured in separate I/O chassis when possible.
	57	Redundant paths to storage devices used by the cluster are required and can be accomplished using either disk mirroring or via LVM's pvlins.
	58	Redundant storage device paths should be configured in separate I/O chassis when possible.
	59	Dual power connected to independent power circuits is recommended.
Heterogeneous Multi System High Availability	60	Cluster configurations can contain a mixture of Superdome and non Superdome nodes.
	61	Care must be taken to configure an even or greater number of nodes outside of the Superdome cabinet
	62	If half the nodes of the cluster are within a Superdome cabinet, a cluster lock is required (4-node maximum cluster size)
	63	If more than half the nodes of a cluster are outside the Superdome cabinet, no cluster lock is required (16-node maximum Serviceguard cluster size).
	64	Up to a 4-node cluster is supported within a single cabinet system (Superdome/ 16-way or Superdome/PA 8800 32-way)
	65	Up to an 8-node cluster is supported within a single cabinet system* (Superdome/ 32-way or Superdome/PA 8800 64-way)
	66	Up to a 16-node cluster is supported within a dual cabinet system* (Superdome/ 64-way or Superdome/PA 8800 128-way)
	67	Cluster lock is required for 2-node configurations
	68	Cluster lock must be powered independently of the cabinet.
	69	Root volume mirrors must be on separate power circuits.
	70	Dual power connected to independent power circuits is highly recommended.

* Superdome 32-way system requires an I/O expansion cabinet for greater than 4 nodes. Superdome 64-way system requires an I/O expansion cabinet for greater than 8 nodes.

NOTE:

"Recommended" refers to configurations that are fully qualified and offer the best bandwidth performance.

"Supported" refers to configurations that are fully qualified, but do not necessarily offer the best performance.

Instant Capacity on Demand (iCOD)

CPU iCOD

Superdome servers can be populated with Itanium 2 CPUs. It is no longer necessary to pay for the additional CPUs until the customer uses them. However with HP's iCOD the remaining 2 CPUs that would cause the cell board to become fully populated can be installed and remain idle. The additional CPUs can be activated instantly with a simple command providing immediate increases in processing power to accommodate application traffic demands.

In the unlikely event that a CPU fails, the HP system will replace the failed CPU on the cell board at no additional charge. The iCOD CPU brings the system back to full performance and capacity levels, reducing downtime and ensuring no degradation in performance.

When additional capacity is required, additional CPUs on a cell board can be brought online. The iCOD CPUs are activated with a single command.

CPU Instant Capacity on Demand (iCOD) can be ordered pre installed on Superdome servers. All cell boards within the

Configuration

Superdome server will be populated with two or four CPUs and the customer orders the number of CPUs that must be activated prior to shipment.

Description	Product Number
Right-to-use dual 1.5 GHz Itanium 2-processor module	A6924A
iCOD right-to-access dual 1.5 GHz Itanium 2-processor module	A6925A

The following applies to CPU iCOD on Superdome servers:

- The number of iCOD processors is selected per partition instead of per system at planning/order time.
- At least one processor per cell in a partition must be a purchased processor.
- Processors are deallocated by iCOD in such a way as to distribute deallocated processors evenly across the cells in a partition. There is no way for a Customer Engineer (CE) or an Account Support Engineer (ASE) or a customer to influence this distribution.
- Reporting for the complex is done on a per partition basis. In other words, all partitions with iCOD processors must be capable of and configured for sending e mail to HP.
- Processors can be allocated and deallocated instantly or after a reboot at the discretion of the user.
- A license key must be obtained prior to either activating or deactivating iCOD processors. A free license key is issued once email connectivity with HP has been successfully established from all partitions with iCOD processors.

Performance Considerations with CPU iCOD:

- Going from one to two to three active CPUs on a cell board gives linear performance improvement
- Going from three to four active CPUs gives linear performance improvement for most applications except some technical applications that push the memory bus bandwidth.
- Number of active CPUs per cell boards should be balanced across partitions. However, minor differences are okay (example: four active CPUs on one cell board and three active CPUs on the second cell board).
- Note that the iCOD software will do CPU activation to minimize differences of number of active CPUs per cell board within a partition.

Cell Board COD

With cell board COD, Superdome servers can be populated with Itanium 2 cell boards (CPU and memory) and it is no longer necessary to pay for the additional cell boards (CPU and memory) until the customer uses them. Additional CPUs and cell boards can be activated instantly with a simple command providing immediate increases in processing power and memory capacity to accommodate application traffic demands.

In the unlikely event that a cell board fails, the HP system will replace the cell board at no additional charge. The COD cell board brings the system back to full performance and capacity levels, reducing downtime and ensuring no degradation in performance.

When additional capacity is required, additional cell boards can be brought online. The COD cell boards are each activated with a single command.

Cell board Capacity on Demand (COD) can be ordered pre installed on Superdome servers. All cell boards within the Superdome server will be populated with two or four CPUs and the customer orders the number of CPUs that must be activated prior to shipment.

iCOD Temporary Capacity

Temporary Capacity for iCOD provides the customer the flexibility to temporarily activate an iCOD processor(s) for a 30-CPU day period. The program includes a temporary Operating Environment (OE) license to use and temporary hardware/software support. The iCOD temporary capacity program enables customers to tap into processing potential for a fraction of the cost of a full activation, to better match expenditures with actual usage requirements and to enjoy the benefits of a true utility model in a capitalized version.

To order iCOD temporary capacity on Superdome, A7067A must be ordered. For more information on iCOD, please refer to the appropriate section in this guide.

Windows Server 2003

Superdome partitions running Windows Server 2003 Datacenter edition (64-bit) do not support CPU iCOD, cell board iCOD and iCOD temporary capacity at this time.

Configuration

Red Hat Enterprise Linux AS 3 and Debian Linux

Superdome partitions running Linux do not support CPU iCOD, cell board iCOD and iCOD temporary capacity.

Utility or Pay-per-Use Program

HP Utility Pricing allows financial decisions on investments to be postponed until sufficient information is available. It allows customers to align their costs with revenues, thereby allowing customers to transition from fixed to variable cost structures. This more flexible approach allows customers to size their compute capacity consistent with incoming revenues and Service Level Objectives. HP Utility Pricing encompasses just-in-time purchased capacity, pay-per-forecast based on planned usage, as well as pay-per-use via metered usage. All offerings are industry leading performance solutions to our customers.

Customers are able to pay for what they use with this new processing paradigm. The usage payments are comprised of both fixed and variable amounts, with the latter based on average monthly CPU usage. Additionally, with HP retaining ownership of the server, technology obsolescence and underutilized processing assets are no longer a customer concern. This is the cornerstone of HP's pay-as-you-go Utility Pricing. Customers will be able to benefit from their servers as a "compute utility". Customers will choose when to apply additional CPU capacity and will only be charged when the additional processing power is utilized. Real-life examples of processing profiles that benefit from Pay per Use are season spikes and month-end financial closings.

The utility program is **mutually exclusive** with iCOD. In order to take part in this program, the utility metering agent (T1322AA) must be ordered.

Windows

Superdome systems running Windows Server 2003 Datacenter edition (64-bit) does not support utility or pay per used program at this time.

Linux

Superdome systems running Linux does not support utility or pay-per-use program.


Upgrades

This section applies to upgrades within HP Integrity Superdome systems. For information on upgrades from existing HP 9000 Superdome systems to HP Integrity Superdome systems, please refer to the upgrade section of the HP 9000 Superdome QuickSpec.

Upgrade Availability

Component	Availability
Add on Upgrades	
Cell Board and CPUs (A6866A and A6924A)	Immediate
2-GB Memory Module (4×512 MB) (A5198A)	Immediate
4-GB Memory Module (4×1 GB DIMMs) (A6863A)	
PCI-X Chassis (A6864A)	Immediate
Superdome I/O Expansion Cabinet (A5861A)	Immediate
Redundant PDCA (A5800A)	Immediate
Model Upgrades	
Superdome 16-way to 32-way	Immediate
Superdome 32-way to 64-way	Immediate
Superdome I/O Expansion Cabinet	Immediate

Upgrade Quick Matrix

Model Upgrade Requirements			
Model Upgrade	Product Number	Configuration Detail Required	Comments
Superdome 16-way to 32-way	A5204A (includes new system backplane)	<ul style="list-style-type: none"> Serial number Existing partition configuration (cell placement, iCOD, OE license type) Media requirements (new DVD, SCSI converter, etc.) 	<ul style="list-style-type: none"> TCE Manager and Deployment Manager High-level design, recommended detailed design Order-entry process through SBW/Watson Config (recommended to identify impact of cell placement, additional software license, etc.) Installation of new backplane included. For additional add on components, see table below.
Superdome 32-way to 64-way	<p>A5202A (right cabinet) must be ordered. A Superdome Service Solution will be required for the upgrade. These solutions (combination of solution level (Foundation, Proactive, Critical) and ordering category (1st system, additional system, additional system later) are selected in SBW/Watson to populate the appropriate product numbers and support. The upgrade includes the right cabinet and integrated cells, CPUs, memory and I/O chassis.</p> <p>NOTE: At least one cell board (A6866A) with one active processor (A6924A), one memory module (A6439A/A6863A), one I/O Card Cage (A6864A) and one Core I/O</p>	<ul style="list-style-type: none"> Serial number Existing partition configuration (cell placement, iCOD, OE license type) Media requirements (new DVD, SCSI converter, etc.) 	<ul style="list-style-type: none"> TCE Manager and Deployment Manager High-level design, detailed design Order-entry process through SBW/Watson Config and Convert to Order required to place cells in correct slot locations Installation included. For additional add on components, see table below.
		DA - 11717 Worldwide — Version 1 — June 30, 2003	Page 27

Upgrades

Card (A5210A) must be ordered when ordering a Superdome 32-way to 64-way upgrade.

Add-on Upgrade Requirements

Component Upgrade	Product Number	Configuration Detail Required	Comments
Memory	A5198A or A6863A	<ul style="list-style-type: none"> Cell count within each partition Memory on each cell board Desired location of ordered memory 	<ul style="list-style-type: none"> No TCE Manager, but optional Deployment Manager Standard order entry process Memory is field installed into existing or new cell boards (included in hw price)
I/O Card Cage	A6864A HA113A1 #588 (installation)	<ul style="list-style-type: none"> Partition info (cell configuration, I/O card cage and I/O cards) 	<ul style="list-style-type: none"> No TCE Manager, but optional Deployment Manager Standard order entry process I/O cage is field installed (requires installation option)
I/O Cards	Depends on I/O card being ordered	<ul style="list-style-type: none"> Validate there are sufficient slots within existing card cage 	<ul style="list-style-type: none"> No TCE Manager, but optional Deployment Manager Standard order entry process I/O cards are field installed into card cages (included in hw price)
Redundant PDCA	A5800A	<ul style="list-style-type: none"> Validate there is an open slot for redundant PDCA 	<ul style="list-style-type: none"> No TCE Manager, but optional Deployment Manager Standard order entry process Redundant PDCA are field installed (included in hw price)
Cell Board into existing partition	A6866A	<ul style="list-style-type: none"> Serial number Existing partition configuration (cell placement, iCOD, OE license type) 	<ul style="list-style-type: none"> No TCE Manager, but optional Deployment Manager Order entry process through Watson Config (recommended to identify impact of cell placement) Installation does not include any partition reconfiguration.
New Partition	A6866A	<ul style="list-style-type: none"> Serial number Existing partition configuration (cell placement, iCOD, OE license type) Media requirements (i.e. new DVD, SCSI converter, etc.) 	<ul style="list-style-type: none"> TCE Manager optional, Deployment Manager recommended High-level design, recommended detailed design Order-entry process through Watson Config (recommended to identify impact of cell placement, additional software license, etc.) Partitions are field installed (included in hw price) Installation does not include any partition reconfiguration

Upgrades

Benefits of Optional Services

Optional Service	Benefits
TCE Manager for adding cells into existing partition or creating a new partition	Enhance TCE by coordinating all of HP's resources focused on fulfilling the customer's solution.
Detailed Design for component add on and Superdome 16-way to 32-way upgrade	Ensures a properly configured solution with the add on or upgrade. Provides pre installation planning for partition reconfiguration and operating environment installation. Minimizes impact to production and reduces customer risk.
Deployment Manager for add-on component	Extend Total Customer Experience by proactively managing/scheduling all field resources required to install and integrate the hardware and to re partition the environment, if required.

Add-on Installation Services

- I/O cards are the ONLY Superdome add on components that are customer installable. However, HP assisted installation is an available option.
- All other Superdome add on components (cell boards, memory, I/O chassis, PDCA) require HP installation. Required installation is available either as a required option or already included in the hardware purchase price.
- Installation options are ordered one per installed product.
- The add-on installation product is as follows:
 - H4725A #588: PCI-X chassis and PCI-X Core I/O installation

Superdome 16-way to 32-way Upgrade

- Includes TCE Manager services.
- Includes Deployment Manager services and installation of the new system backplane.
- Installation for any additional items into the cabinet should be ordered using the add-on process.
- Does not include any partition reconfiguration or new partition creation/OS load.
- Does not require site readiness or preparation

Superdome 32-way to 64-way Upgrade

- Includes TCE Manager and Deployment Manager services.
- Includes:
 - Detailed Design
 - Deployment Manager services
 - Site Environmental Services
 - Right cabinet
- Factory installation of internal components (cell boards, memory, PCI X chassis, I/O cards and redundant PDCA) into the right cabinet
- Items for the left cabinet should be ordered using the add on installation options.
- Does not include partition reconfiguration or new partition creation/OS load.

Superdome 32-way to 64-way Minimum Order Requirements

The requirements below are in place to facilitate delivering the right cabinet with a high level of product quality:

- Quantity 1 cell board (1 active CPU minimum)
- Quantity 1 memory module (2 GB, set of 4 512-MB DIMMs)
- Quantity 1 PCI X chassis
- Quantity 1 core I/O card

iCOD for Add-on and Upgrades

- iCOD may be ordered on add-on cell boards, Superdome 16-way to 32-way upgrade and Superdome 32-way to 64-way upgrade.
- At least one CPU must be active (purchased) per cell boards.
- The iCOD client agent license should be ordered, B9073AA.

Upgrades

iCOD for Add-on and Upgrades

- The customer's current system configuration is required for all add ons (except I/O cards) and upgrades to ensure an accurate design and proper installation:
 - For CSS and BCS customers, the on site ASE collects this data.
 - Foundation customers may choose to have HP collect the data or provide it themselves. Customer viewable documentation on extraction of this data will be available.
- The Solution Architect uses the current configuration information to accurately document the placement of component add on and model upgrades. Ultimately, this information is provided to the installation CE.
- New Watson Configurator/SBW functionality allows for an easy process to upgrade from the original system configuration.
 - The current configuration data is entered into the tool, the "Upgrade" option select and only the add on component or model upgrade products are quoted.
- Watson Configuration/SBW tools are required as follows:
 - For an add-on component, not required to use Watson Configurator/SBW to create configuration, but highly recommended to identify optimal component placement
 - For a Superdome 16-way to 32-way upgrade, not required to use Watson Configurator/SBW to create configuration, but highly recommended to identify optimal component placement
 - For a Superdome 32-way to 64-way upgrade, required to use Watson Configurator/SBW and Convert to Order to enable correct placement of cell boards and I/O within right cabinet.

Memory

Total Amount Memory per Cell	No. of 512 MB	No. of 1 GB	E0	E1	E2	E3	E4	E5	E6	E7	E8	E9	EA	EB	EC	ED	EE	EF
			0A- 0B	1A- 1B	2A- 2B	3A-3B	4A- 4B	5A- 5B	6A- 6B	7A- 7B	8A- 8B	9A-9B	AA- AB	BA-BB	CA- CB	DA- DB	EA-EB	FA-FB
4 GB	8		512 MB	512 MB	512 MB	512 MB												
8 GB		8	1 GB	1 GB	1 GB	1 GB												
8 GB	16		512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB								
16 GB		16	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB								
16 GB	32		512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB
32 GB		32	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB
12 GB	8	8	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB								
24 GB	16	16	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB
28 GB	8	24	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB
2 GB	4		512 MB	512 MB														
4 GB		4	1 GB	1 GB														
12 GB	24		512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB	512 MB				
24 GB		24	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB				
20 GB	8	16	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	1 GB	512 MB	512 MB	512 MB	512 MB				

Ex Echelon number, i.e. Echelon 0 consists of 2 DIMMs, 1 on A side and 1 on B side.

0A 0B refers to two DIMMs in Echelon 0, A and B side.

Supported List of Mixed DIMM Sizes in Superdome

Total Amount of Memory Per Cell (GB)	Number of 512 MB DIMMs	Number of 1 GB DIMMs
2	4	0
4	0	4
12	24	0
20	8	16
24	0	24

Supported List of Mixed DIMM Sizes in Superdome with Best Performance (Recommended)

Total Amount of Memory Per Cell (GB)	Number of 512 MB DIMMs	Number of 1 GB DIMMs
4	8	0
8	0	8
8	16	0
12	8	16
16	0	0
16	32	16
24	16	16
28	8	24
32	0	32

TechSpecs

Superdome Specifications	SPU Model Number	Superdome 16-way	Superdome 32-way	Superdome 64-way
	SPU Product Number	A6113A	A5201A	A5201A+A5202A
	TPC-C disclosure (HP-UX)	N/A	N/A	TBD
	TPC-C disclosure (Windows)	N/A	N/A	707,102 tpmC (Windows Server 2003 Datacenter Edition with SQL Server 2000 (64-bit version))
	Number of CPUs	2-16	2-32	6-64
	Itanium 2 Processor	1.5 GHz, 6 MB cache	1.5 GHz, 6 MB cache	1.5 GHz, 6 MB cache
	Memory (with 512 MB or 1 GB DIMMs)	2-128 GB	2-256 GB	6-512 GB
	2-way or 4-way Cells	1-4	1-8	3-16
	12-slot PCI-X I/O chassis	1-4	1-8	1-16
	NOTE: SPU cabinet must be filled first before placing I/O chassis in I/O expansion cabinet	No I/O expansion cabinet required.	I/O expansion cabinet required if number of I/O chassis is greater than 4.	I/O expansion cabinet required if number of I/O chassis is greater than 8. A second I/O expansion cabinet is required if the number of I/O chassis is greater than 14.
	Number of Partitions without I/O expansion cabinet	1-4	1-4	1-8
	Number of Partitions with I/O expansion cabinet	N/A	1-8	1-16
	HP-UX revision	HP-UX 11i version 2: August 2003		
	Windows revision	Windows Server 2003, Datacenter edition for Itanium 2; 2 4 months after HP-UX ship release		
	Linux revision	Red Hat Enterprise Linux AS 3: December 2003; Debian Linux: TBD		
	RS-232C Serial Ports	Y	Y	Y
	10/100Base-T Ethernet	Y	Y	Y
	DIMM Density (MB)	512/1024	512/1024	512/1024
	Site planning and installation included	Y	Y	Y
	Maximum Heat dissipation (BTUs/hour)	28,969	41,614	83,288
	Typical Heat dissipation (BTUs/hour)	20,131	33,439	66,877
	Depth (in/mm)	48.03/1,220	48.03/1,220	48.03/1,220
	Width (in/mm)	30/762	30/762	60/1,524
	Height (in/mm)	77.16/1,960	77.16/1,960	77.16/1,960
	Weight (lbs/kg)	1102.31/500	1318.36/598	2636.73/1,196
	Electrical Characteristics			
	AC input power-Option 7: 3 phase 5 wire input	200 240 VAC phase to neutral, 5 wire, 50/60 Hz		
	AC input power-Option 6: 3 phase 4 wire input	200 240 VAC phase to phase, 4 wire, 50/60 Hz		
	Current requirements at 220V-240V			
	Option 7:	24 A	24 A	24 A

TechSpecs

3-phase 5-wire input			
Option 6:	44 A	44 A	44 A
3-phase 4-wire input			
Required Power Receptacle- Options 6 and 7	None. Cord, plug and included. Receptacle should be ordered separately. Electrician must hard wire receptacle to site power.		
Maximum Input Power (watts)	8,490	12,196	24,392
Typical Input Power (watts)	5,900	9,800	19,600
	4 cells, 32 GB, 4 I/O chassis with 6 PCI each	8 cells, 32 GB, 4 I/O chassis with 6 PCI each	16 cells, 32 GB, 4 I/O chassis with 6 PCI each

Environmental Characteristics

Acoustics	65 dB
Operating temperature	68° to 86°F (20°C to 30°C)
Non-operating temperature	-40° to 158°F (-40° to 70°C)
Maximum rate of temperature change	68°F/hr (20°C/hr)
Operating relative humidity	15% to 80% @ 86°F (30°C)
Operating altitude	0 to 10,000 ft (0 to 3.1 km)
Non-operating altitude	0 to 15,000 ft (0 to 4.6 km)

Regulatory Compliance

Safety	IEC 950:1991+A1, +A2, +A3, +A4; EN60950:1992+A1, +A2, +A3, +A4, +A11; UL 1950, 3rd edition; cUL CSA C22.2 No. 950 95
--------	---

Key Dates

First CPL date	6/03
First ship date	3Q03

Dimensions

Height	5.25 ft or 6.43 ft (1.6 meters or 1.96 meters)
Depth	45.5 in (115.67 cm)(same depth as 32W)
Width	24.0 in (60.96 cm)

Electrical Characteristics

AC input power	200-240 VAC, 50/60 Hz
Current requirements at 200V240V	16 A
Typical maximum power dissipation (watts)	2290
Maximum power dissipation (watts)	3200
Environmental Characteristics	Same as Superdome

TechSpecs

Superdome I/O Expansion (IOX) Cabinet Specifications	Maximum Number of I/O Chassis Enclosures (ICEs)*	3
	Peripherals Supported	All peripherals qualified for use with Superdome and/or for use in a Rack System E are supported in the I/O expansion cabinet as long as there is available space. Peripherals not connected to or associated with the Superdome system to which the I/O expansion cabinet is attached may be installed in the I/O expansion cabinet.
	Servers Supported	No servers except those required for Superdome system or High Availability Observatory or ISEE may be installed in an I/O expansion cabinet.
	Superdome Models Supported	Superdome 32-way Superdome 64-way
	Relevant Product Numbers	12 slot PCI X Chassis for Rack System E Expansion Cabinet A6864AZ I/O expansion cabinet Power and Utilities Subsystem A5861A I/O Chassis Enclosure for 12 slot PCI X Chassis A5862A

* Each ICE holds two I/O card cages or 24 PCI-X I/O slots.

APC SL20KFB2 Specifications	Description	APC Silcon, 20000VA/20000W, Input 115/200 3PH, 120/208 3PH, 127/220 3PHV/ Output 115/200 3PH, 120/208 3PH, 127/220 3PHV, Interface Port DB-25 RS--232, Contact Closure	
	General Features	0% to 95% non-condensing, 200% overload capability, Audible Alarms, Built in static bypass switch, Delta Conversion On line Technology, Environmental Protection, Event logging, Extendable Run Time, Full rated output available in kW, Input Power Factor Correction, Intelligent Battery Management, LCD Alphanumeric Display, Overload Indicator, Paralleling Capability, Sine wave output, SmartSlot, Software, Web Management	
	Includes	Parallel Card, Triple Chassis for three SmartSlots, User Manual, Web/SNMP Management Card	
	Spare parts kits	See APC website www.apcc.com	
	Documentation	User Manual and Installation Guide	
	Input	Nominal input voltage	115/200 3PH, 120/208 3PH, 127/220 3PH V
		Input frequency	50 Hz programmable +/- 0.5, 1, 2, 4, 6, 8%; 60 Hz programmable +/- 0.5, 1, 2, 4, 6, 8%
		Input connection type	Hardwire 5-wire (3PH + N + G)
		Input voltage range for main operations	170-230 (200 V), 177-239 (208 V), 187-242 (220 V) V
	Batteries	Typical backup time at half load	36.7 minutes
		Typical backup time at full load	10.7 minutes
		Battery type	Maintenance-free sealed Lead-Acid battery with suspended electrolyte: leak proof
		Typical recharge time **	2 hours
	Physical	Maximum height dimensions	55.12 in (140.00 cm)
		Maximum width dimensions	39.37 in (100.00 cm)
		Maximum depth dimensions	31.50 in (80.01 cm)
		Net weight	1,290.00 lbs (586.36 kg)
		Shipping Weight	1,340.00 lbs (609.09 kg)
		Shipping Height	66.93 in (170.00 cm)
		Shipping Width	43.31 in (110.00 cm)
		Shipping Depth	35.43 in (90.00 cm)
		Color	Dark green (NCS 7020 B50G), Light gray (NCS 2703

TechSpecs

Communications and Management	Units per Pallet	G84Y) 1.0
	Interface port	DB-25 RS-232, Contact Closure
	Smart Slot Interface Quantity	2
	Pre-Installed SmartSlot Cards	AP9606
	Control panel	Multi-function LCD status and control console
	Audible alarm	Beep for each 52 alarm conditions
	Emergency Power Off (EPO)	Yes
Environmental	Optional Management Device	See APC website www.apcc.com
	Operating Environment	32° to 104°F (0° to 40 °C)
	Operating Relative Humidity	0% to 95%
	Operating Elevation	0 to 3333 ft (0 to 999.9 m)
	Storage Temperature	-58° to 104°F (-50° to 40 °C)
	Storage Relative Humidity	0% to 95%
	Storage Elevation	0 to 50,000 ft (0 to 15,000 m)
Conformance	Audible noise at 1 meter from surface of unit	55 dBA
	Online thermal dissipation	4,094 BTU/hour
	Protection Class	NEMA 1, NEMA 12
	Approvals	EN 55022 Class A, ISO 9001, ISO 14001, UL 1778, UL Listed, cUL Listed
	Standard warranty	One-year repair or replace, optional on-site warranties available, optional extended warranties available
	Optional New Service	See APC website www.apcc.com

TechSpecs

Superdome Supported I/O

Card	Product #	HP-UX 11i v2	Windows Server 2003 Datacenter Edition	Linux
LAN/WAN				
FDDI Universal PCI Adapter	A3739B	Yes	No	No
1000Base-SX PCI LAN Adapter	A4926A	Yes - March 04	No	No
1000Base-T PCI Gigabit Ethernet LAN Adapter	A4929A	Yes - March 04	No	No
PCI 10/100Base-T LAN Adapter	A5230A	Yes	No	No
PCI 4-port 100Base-TX LAN Adapter	A5506B	Yes	No	TBD
PCI ATM 155 Mbps MMF Adapter	A5513A	Yes	No	No
PCI Token Ring 4/16/100 Hardware Adapter	A5783A	Yes	No	No
PCI 2-port 100Base T 2-Port Ultra2 SCSI	A5838A	Yes - No boot or Serviceguard support	No	No
PCI 1000Base-T Gigabit Ethernet Adapter	A6825A	Yes	No	No
Windows/Linux PCI 1000Base-T Gigabit Ethernet Adapter (Copper)	A7061A	No	Yes	TBD
Windows/Linux PCI 1000Base-SX Gigabit Ethernet Adapter (Fiber)	A7073A	No	Yes	TBD
PCI 1000Base-SX Gigabit Ethernet Adapter	A6847A	Yes	No	No
SCSI				
PCI Ultra160 SCSI Adapter	A6828A	Yes	No	No
Windows/Linux Ultra160 SCSI Adapter	A7059A	No	Yes	TBD
Windows/Linux Dual-channel Ultra160 SCSI Adapter	A7060A	No	Yes	TBD
PCI Dual-channel Ultra160 SCSI Adapter	A6829A	Yes	No	No
RAID				
PCI-X RAID Smart Array 6402 U320, 2CH	A9890A	No	Yes	Yes
PCI-X RAID Smart Array 6404 U320, 4CH	A9891A	No	Yes	Yes
FC				
PCI 2X Fibre Channel Adapter	A5158A	Yes - No boot support	No	No
PCI 2-Gb Fibre Channel Adapter	A6795A	Yes	No	No
PCI X 2-Gb FCA2404 Fibre Channel HBA	AB232A	No	Yes - No boot support	No
PCI X 2-Gb FCA2214 Fibre Channel HBA	AB234A	No	No	TBD
Miscellaneous				
PCI HyperFabric2 fiber adapter	A6386A	Yes	No	No
PCI 8-port serial MUX adapter	A6748A	Yes	No	No
PCI 64-port serial MUX adapter	A6749A	Yes	No	No
Dual-port PSI Adapter	J3525A	Yes	No	No

TechSpecs

Superdome Supported Online Storage

Storage Device	HP-UX 11i v2	Windows Server 2003 Datacenter Edition	Linux
XP 48/512	Yes	Yes	Yes
XP128/1024	Yes	Yes	Yes
VA7100	Yes	No	Yes
VA7400	Yes	No	Yes
VA 7410/7110	Yes	Yes	Yes
MSA1000	Yes - 2H 2004	Yes	Yes
EVA 5000	Yes	Yes	Yes
EVA 3000	Yes	Yes	Yes
StorageWorks 4400 series	No	Yes	No
FC10	Yes	No	No
SC10	Yes	No	No
DS2100	Yes	No	Yes
DS2110	Yes	No	Yes
DS2300	Yes	No	Yes
DS2405	Yes	No	No
EMC Symmetrix 3000	Yes	TBD	No
EMC Symmetrix 5000	Yes	TBD	No
EMC Symmetrix 5500	Yes	TBD	No
EMC Symmetrix 8000	Yes	TBD	No
EMC DMX Series	Yes	TBD	TBD
SAN 2/8	Yes	No	Yes
SAN 2/8 EL	Yes	No	Yes
SAN 2/16	Yes	No	Yes
SAN 2/16 EL	Yes	No	Yes
StorageWorks Core 2/64	Yes	Yes	Yes
StorageWorks Edge 2/24	Yes	No	Yes
StorageWorks Edge 2/32	Yes	No	Yes
StorageWorks SAN Director 2/64	Yes	Yes	Yes
StorageWorks SAN Director 2/140	Yes	No	Yes

TechSpecs

Superdome Supported Nearline Storage

Storage Device	HP-UX 11i v2	Windows Server 2003 Datacenter Edition	Linux
ESL9595 with SDLT 220 and 320	Yes	No	Yes
ESL9595 with Ultrium 230 and 460 drives	Yes	No	Yes
ESL9322 with SDLT 220 and 320	Yes	No	Yes
ESL9322 with Ultrium 230 and 460 drives	Yes	No	Yes
MSL5000 series with Ultrium 230 drives	Yes	No	Yes
MSL5000 series with SDLT 220 drives	Yes	No	Yes
MSL5000 series with SDLT 320 drives	Yes	No	Yes
MSL6000 series with Ultrium 460 drives	Yes	No	Yes
SSL1016 with DLT1	Yes	No	Yes
SSL1016 with SDLT 320	Yes	No	Yes
SSL1016 with Ultrium 460	Yes	No	Yes
Tape Autoloader 1/8	Yes	No	Yes
NSR 1200 FC/SCSI router for MSL series libraries	Yes	No	Yes
NSR e1200, e1200-160 FC/SCSI router for MSL libraries	Yes	No	Yes
NSR e2400, e2400-160 FC/SCSI router for ESL libraries	Yes	No	Yes
NSR 2402 FC/SCSI router for ESL series libraries	Yes	No	Yes
Optical Jukebox 2200mx	Yes	No	No
Optical Jukebox 1200mx	Yes	No	No
Optical Jukebox 700mx	Yes	No	No
Optical Jukebox 600mx	Yes	No	No
Optical Jukebox 300mx	Yes	No	No
Optical Jukebox 220mx	Yes	No	No
Optical Jukebox 9100mx	Yes	No	No
Ultrium 460 Standalone/Rack	Yes	No	Yes
Ultrium 230 Standalone/Rack	Yes	No	Yes
Ultrium 215 Standalone/Rack	Yes	No	Yes
DVD ROM Rack	Yes	No	Yes
TA5300 Tape Array	Yes	No	Yes
DDS-4 Standalone/Rack	Yes	No	Yes
DDS-4×6 Standalone	Yes	No	Yes
DDS-5 Standalone/Rack	Yes	No	Yes
DLT-80 Standalone/Rack	Yes	No	Yes
DLTVS80 Standalone/Rack	Yes	No	Yes

- **NOTES:**

All shipments of SCSI devices for Superdome except HVD10 and SC10 are supported with standard cables and auto termination enabled. Only the Surestore Disk System HVD10 (A5616AZ) and the Surestore Disk System SC10 (A5272AZ) will use disabled auto termination and In-Line Terminator cables.

- Each A5838A PCI 2-port 100Base-T 2-port Ultra2 SCSI card that supports a Surestore Disk System SC10 (A5272AZ) will need quantity two (2) of product number C2370A (terminator); otherwise it must have a terminated cable in place prior to HP UX boot.

TechSpecs

© Copyright 2003 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice.

Microsoft and Windows Server 2003 are US registered trademarks of Microsoft Corporation. Intel and Itanium are US registered trademarks of Intel Corporation.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.